Digital Photography

ray optics, aperture, depth of field, exposure, sensor noise, demosaicking, denoising

CSC420
David Lindell
University of Toronto
cs.toronto.edu/~lindell/teaching/420
Slide credit: Babak Taati ←Ahmed Ashraf ←Sanja Fidler

TCIG
Toronto Computational Imaging Group

Announcements

- HW 2 is out (due Oct 16)

Let’s say we have a sensor…

digital sensor (CCD or CMOS)

… and an object we like to photograph

real-world object

digital sensor (CCD or CMOS)

What would an image taken like this look like?

Bare-sensor imaging

real-world object

digital sensor (CCD or CMOS)

Bare-sensor imaging

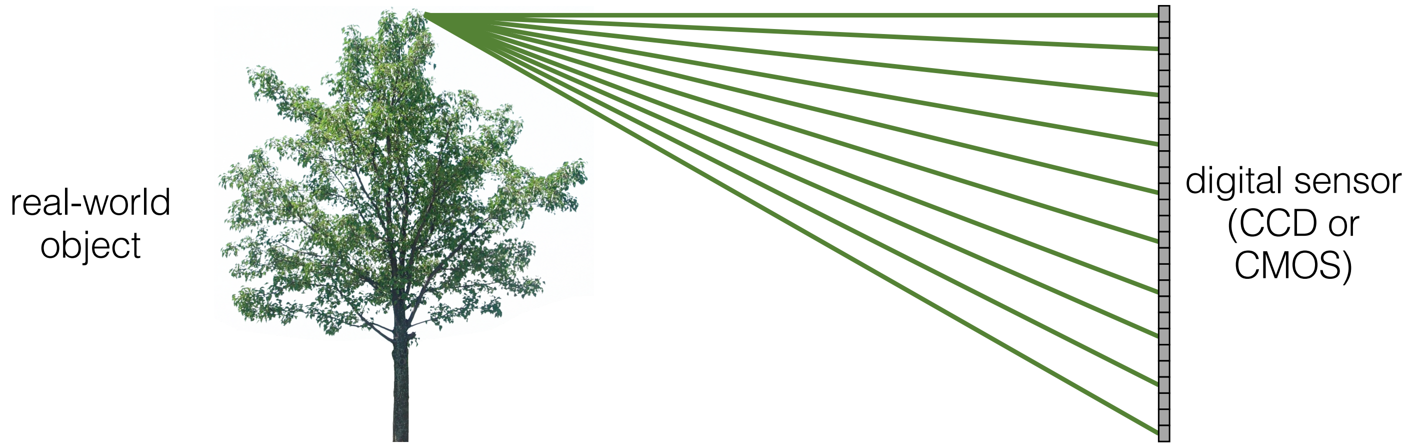

Bare-sensor imaging

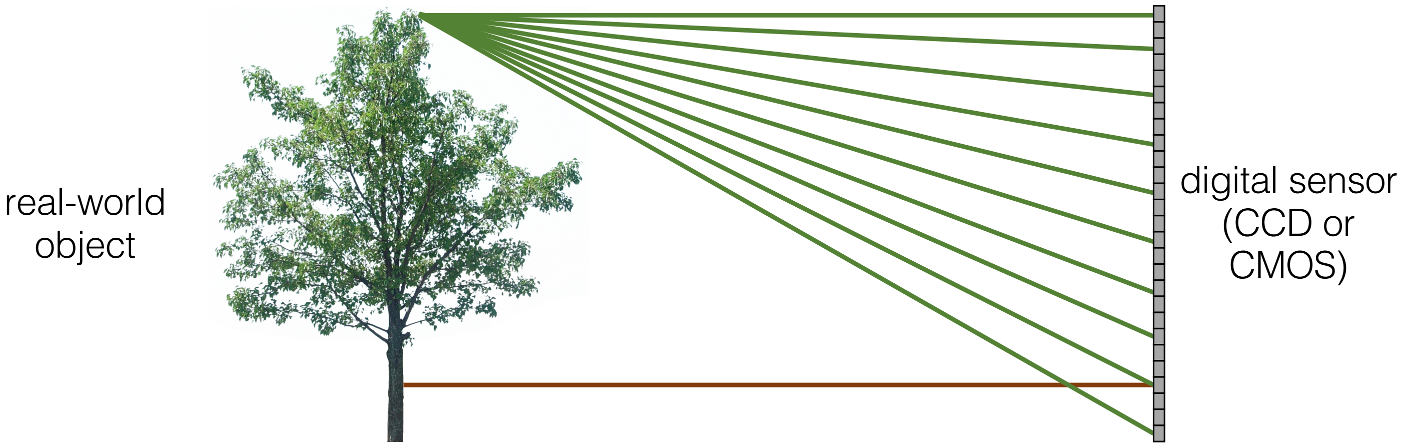

Bare-sensor imaging

real-world object

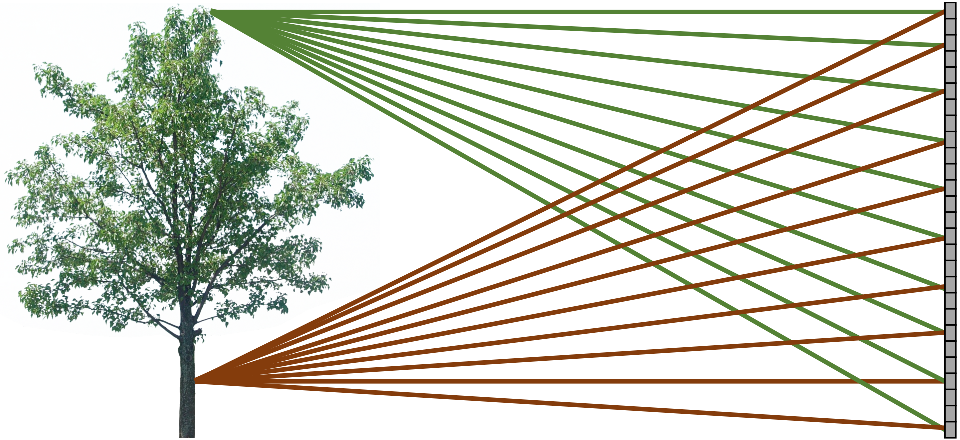

digital sensor (CCD or CMOS)

All scene points contribute to all sensor pixels

What does the image on the sensor look like?

Bare-sensor imaging

All scene points contribute to all sensor pixels

What can we do to make our image look better?

real-world
object

digital sensor
(CCD or
CMOS)

Let’s add something to this scene

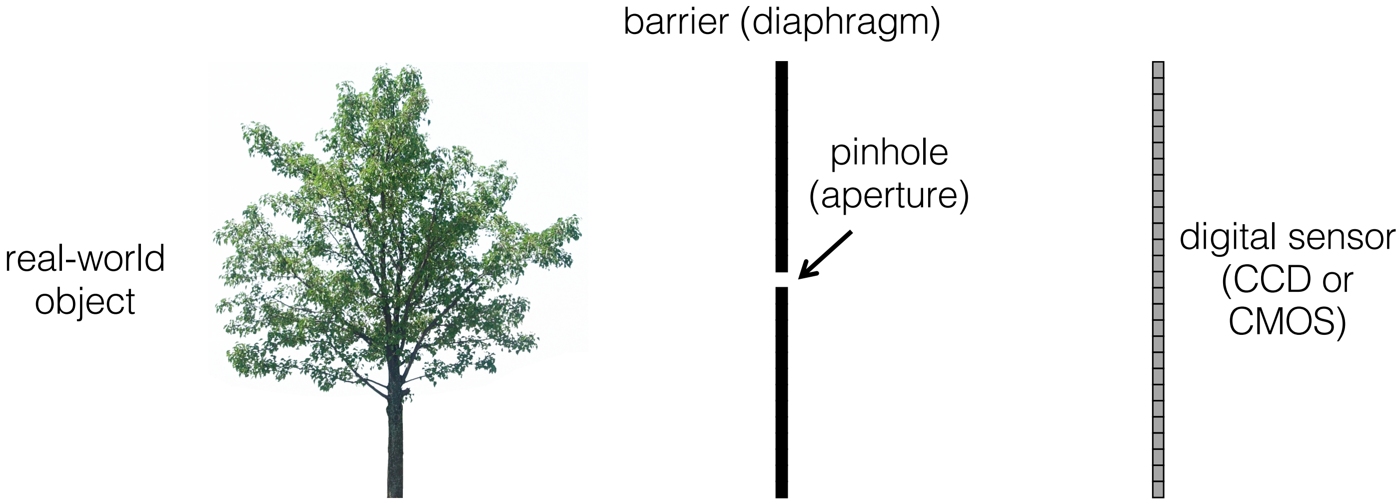

What would an image taken like this look like?

Pinhole imaging

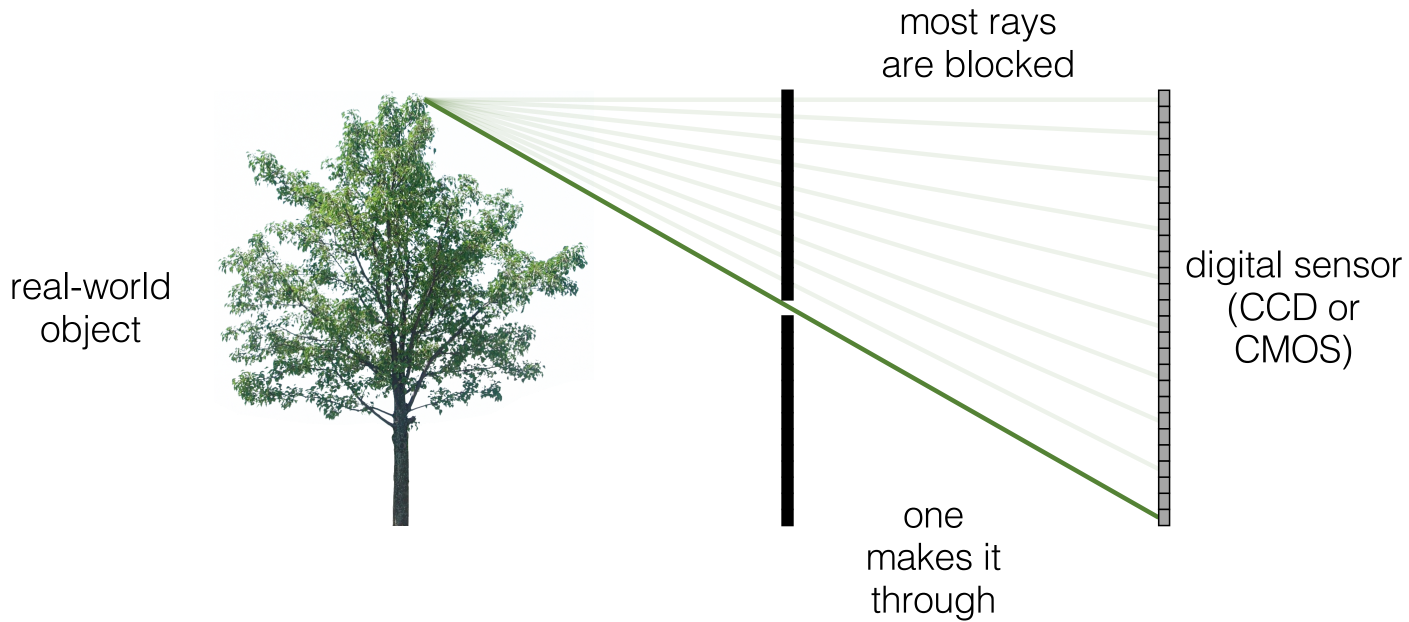

Pinhole imaging

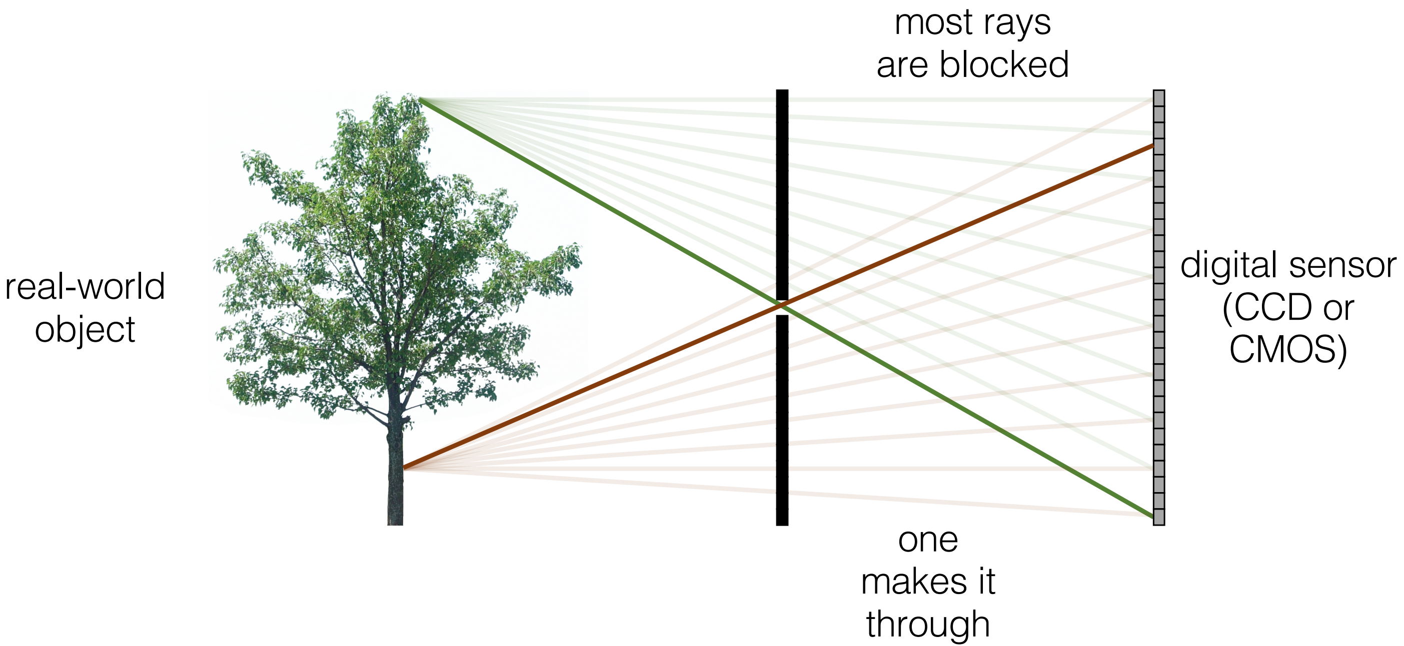

Pinhole imaging

real-world object

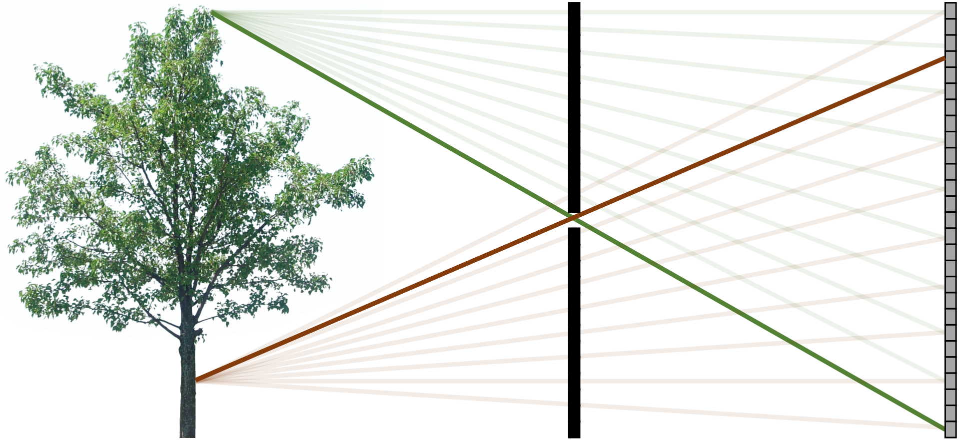

digital sensor (CCD or CMOS)

Each scene point contributes to only one sensor pixel

What does the image on the sensor look like?

Pinhole imaging

real-world object

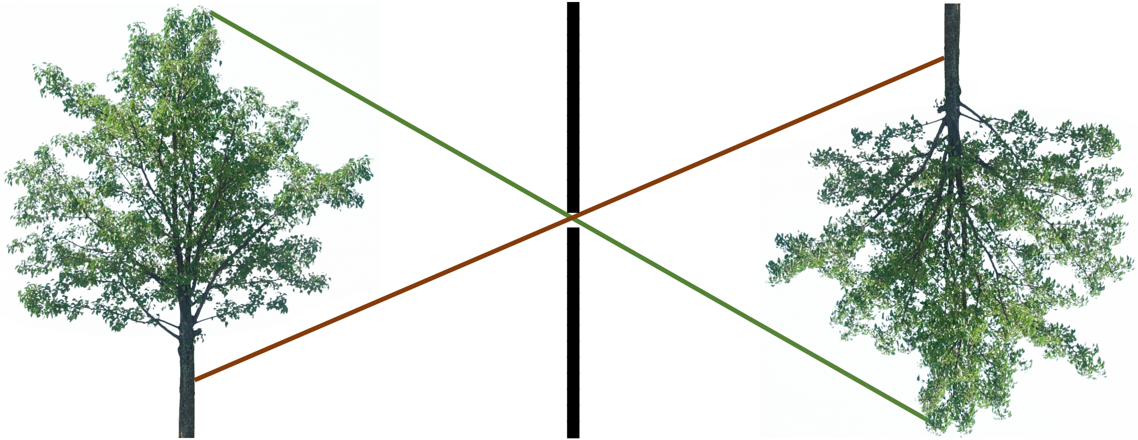

copy of real-world object (inverted and scaled)

Pinhole camera terms

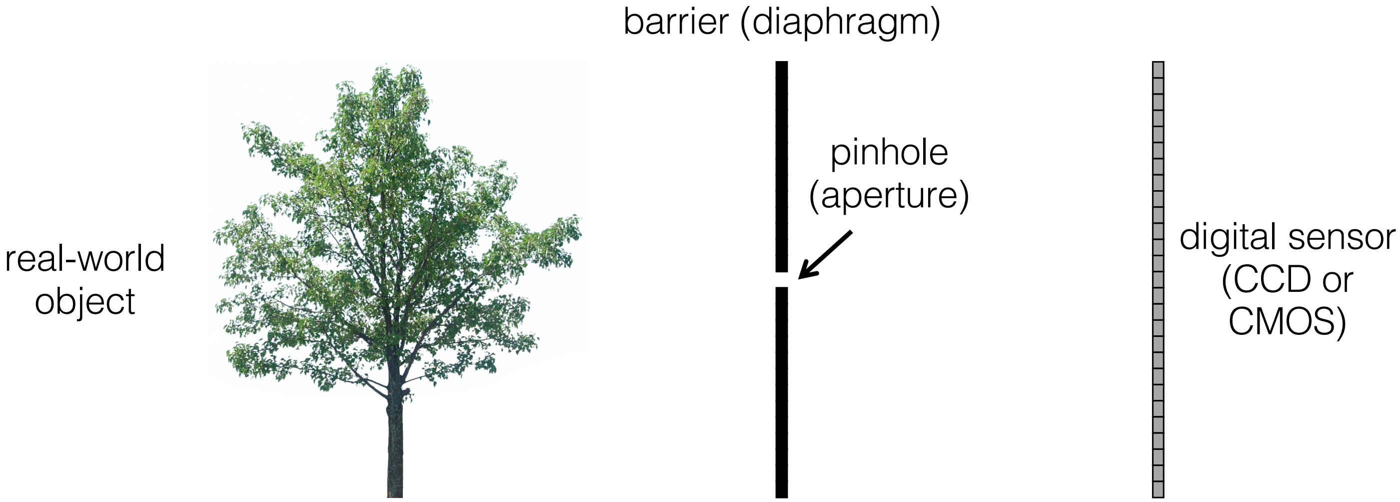

Pinhole camera terms

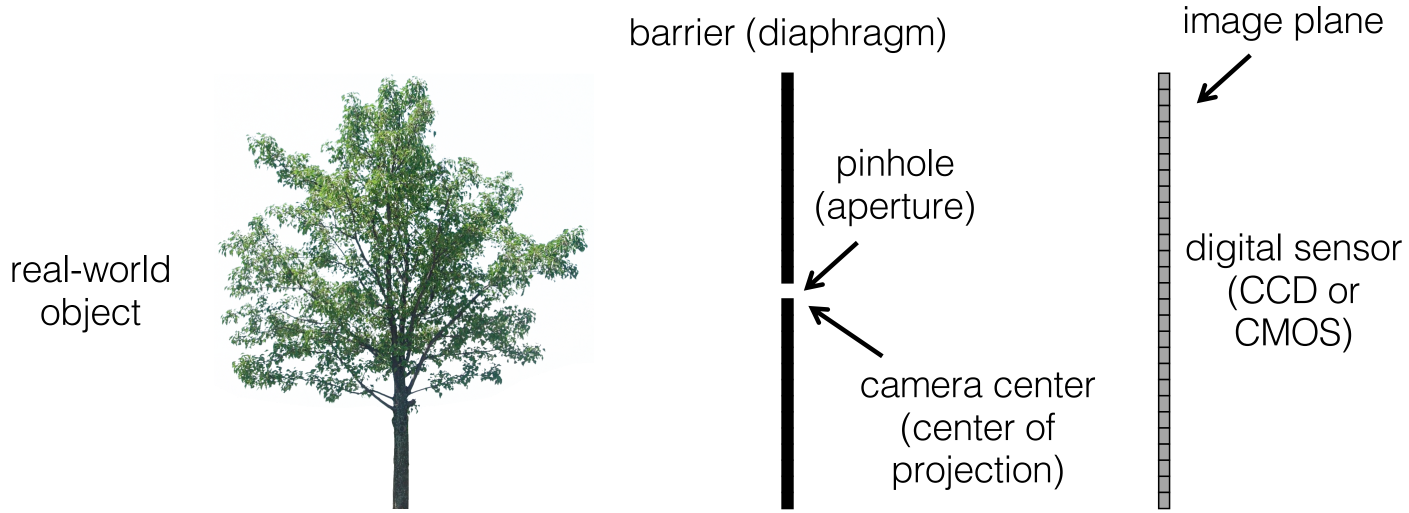

Focal length

real-world
object

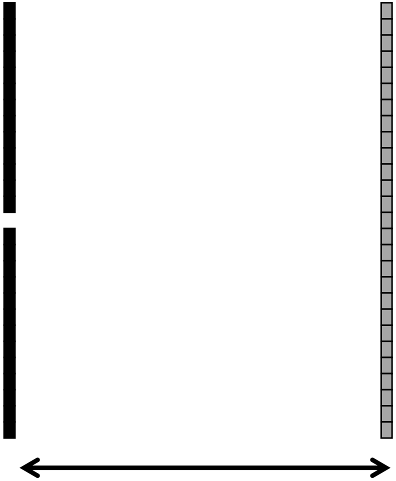

focal length f

Focal length

What happens as we change the focal length?

real-world object

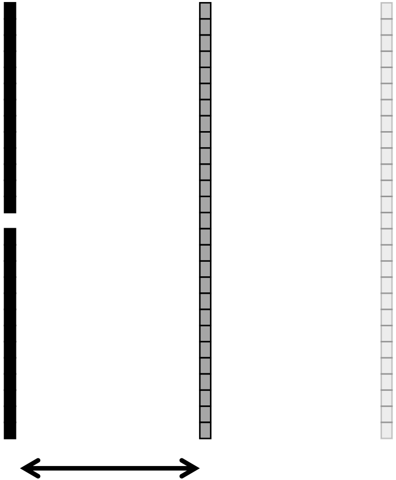

focal length 0.5 f

Focal length

What happens as we change the focal length?

real-world object

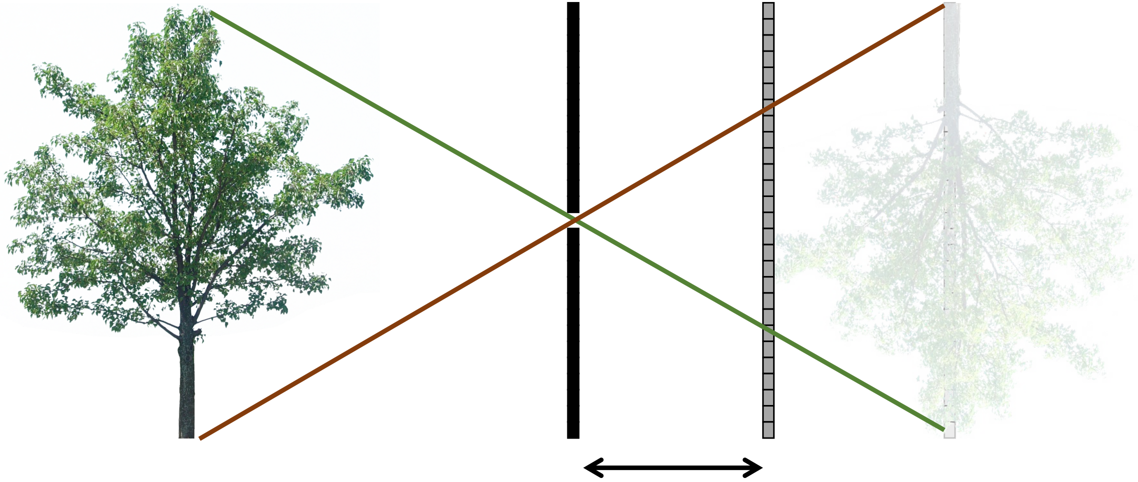

focal length 0.5 f

Focal length

What happens as we change the focal length?

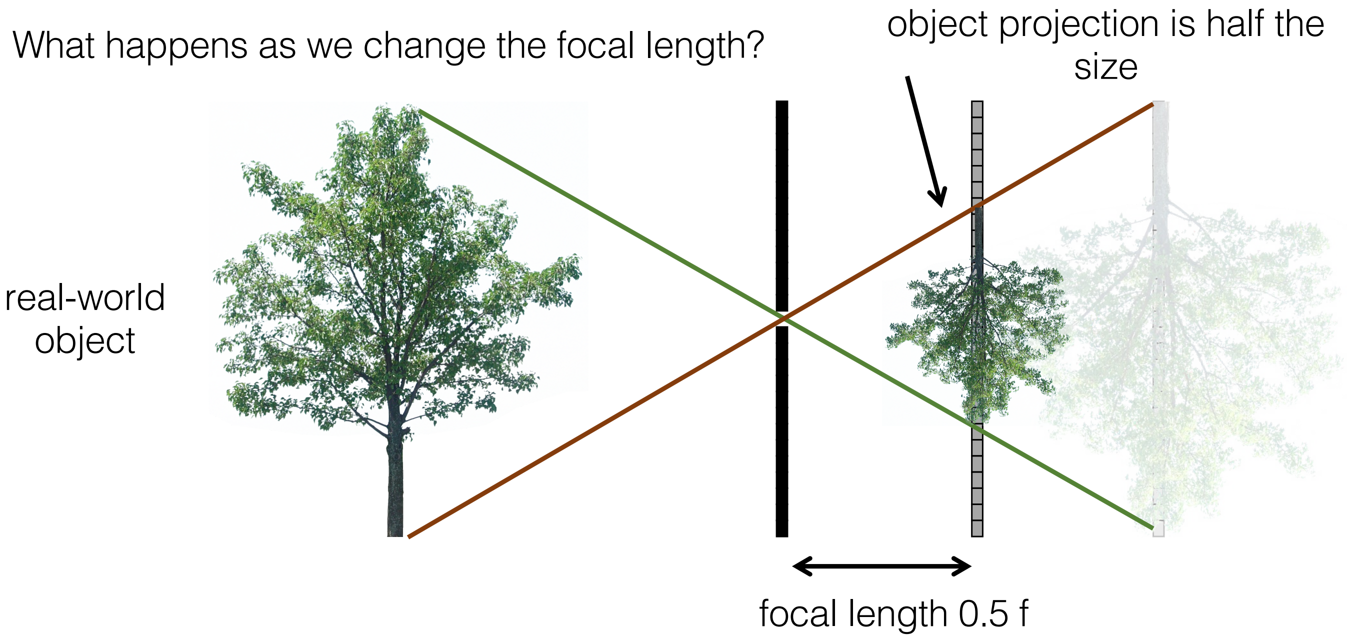

Pinhole size

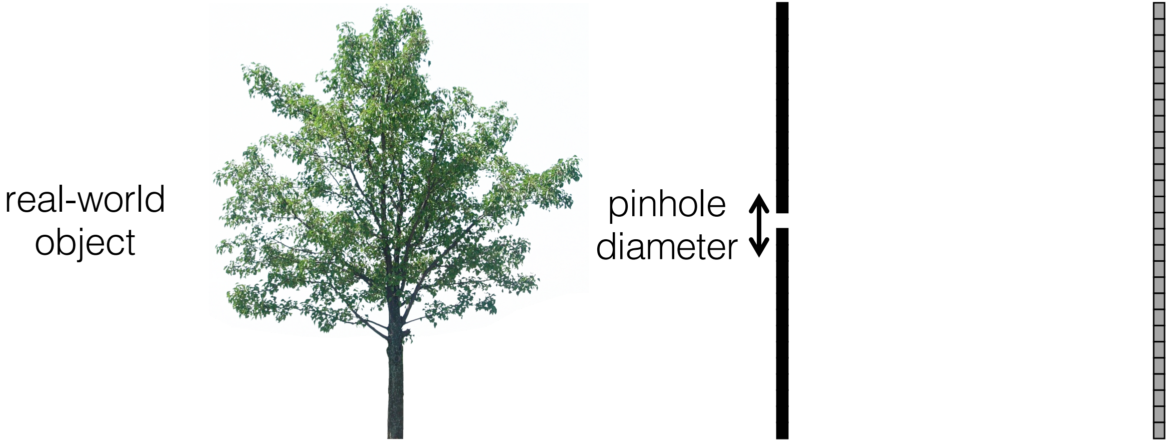

Ideal pinhole has infinitesimally small size

- In practice that is impossible.

Pinhole size

What happens as we change the pinhole diameter?

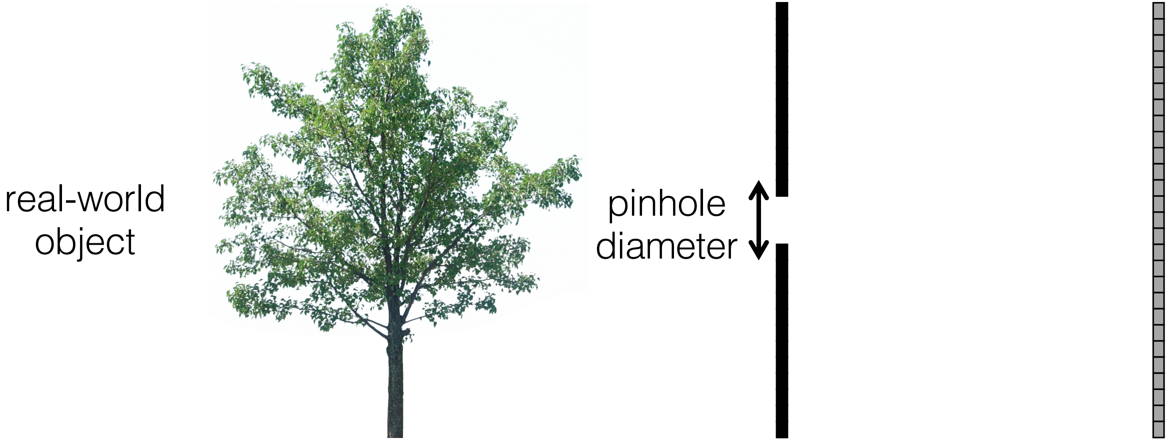

Pinhole size

What happens as we change the pinhole diameter?

real-world object

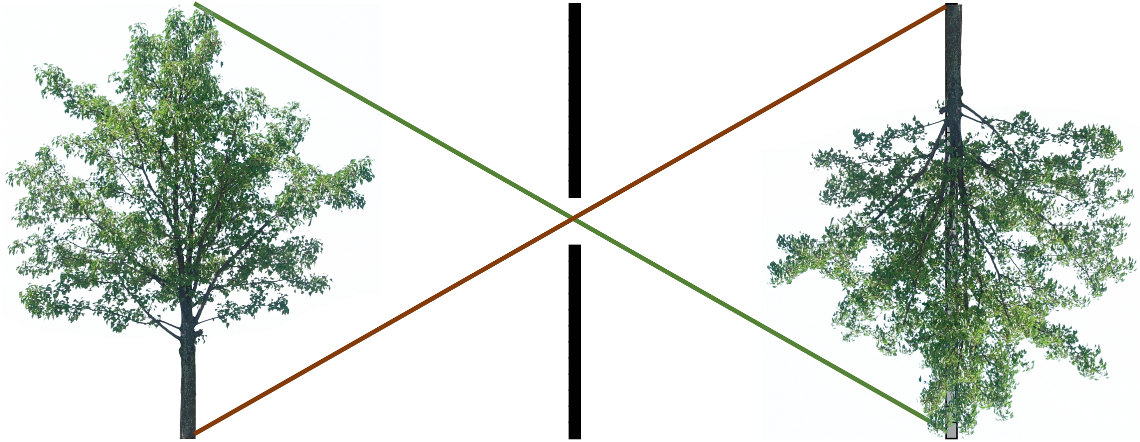

Pinhole size

What happens as we change the pinhole diameter?

real-world object

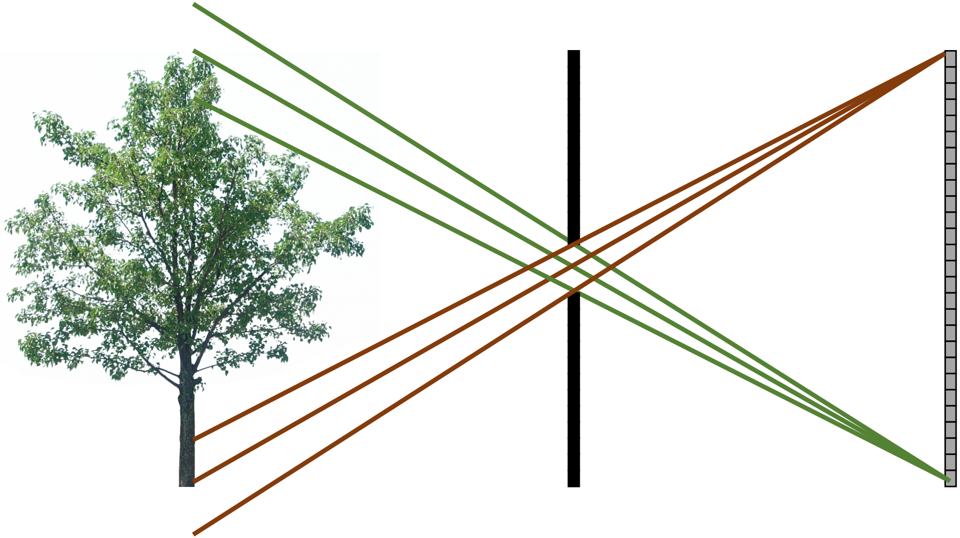

Pinhole size

What happens as we change the pinhole diameter?

object projection becomes blurrier

real-world object

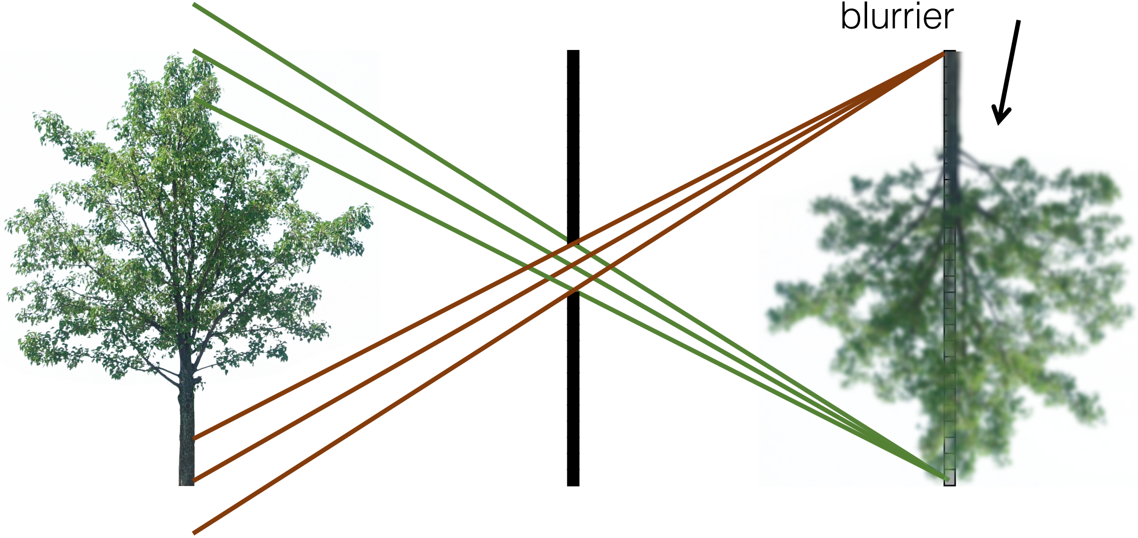

Pinhole size

What happens as we change the pinhole diameter?

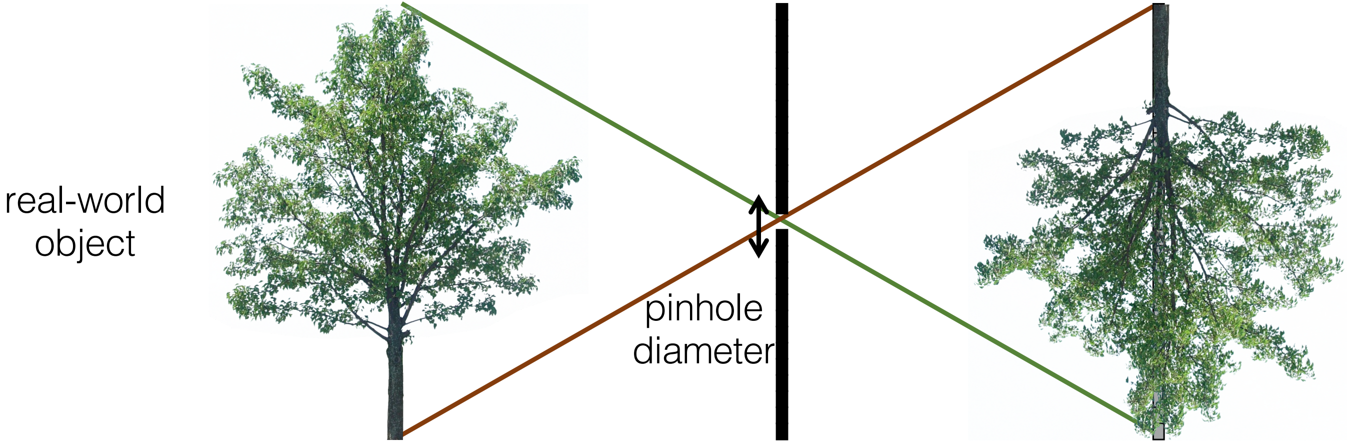

Will the image keep getting sharper the smaller we make the pinhole?

Diffraction limit

A consequence of the wave nature of light

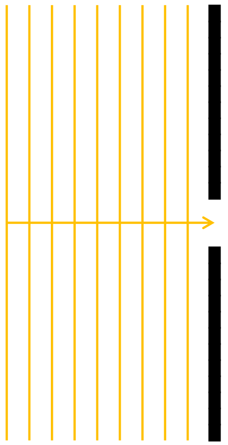

What do geometric optics predict will happen?

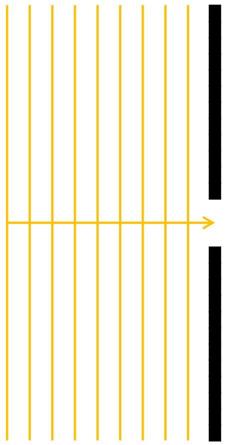

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

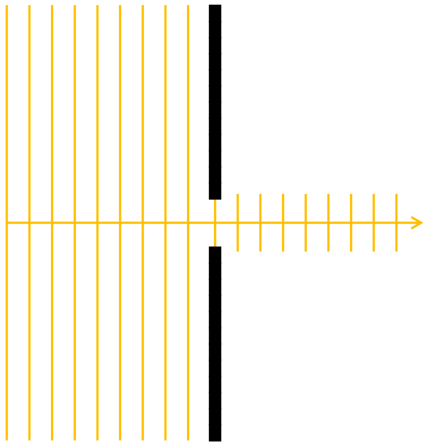

What do geometric optics predict will happen?

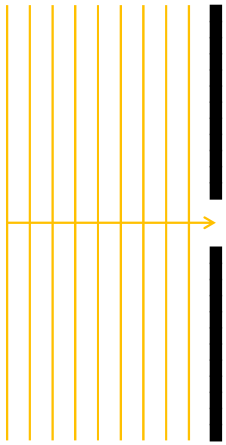

What do wave optics predict will happen?

Diffraction limit

A consequence of the wave nature of light

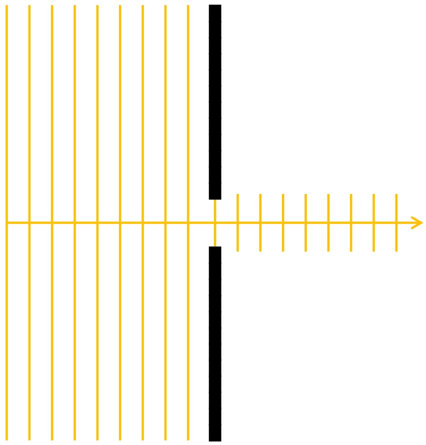

What do geometric optics predict will happen?

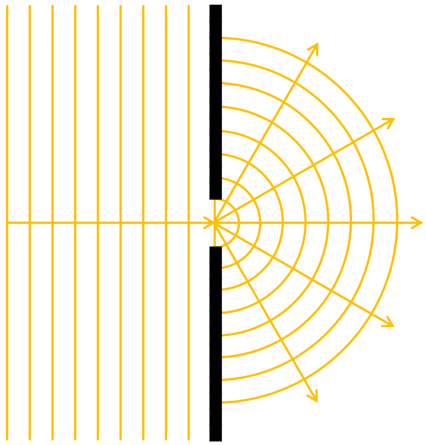

What do wave optics predict will happen?

Diffraction limit

Diffraction pattern = Fourier transform of the pinhole.

- Smaller pinhole means bigger Fourier spectrum.
- Smaller pinhole means more diffraction.

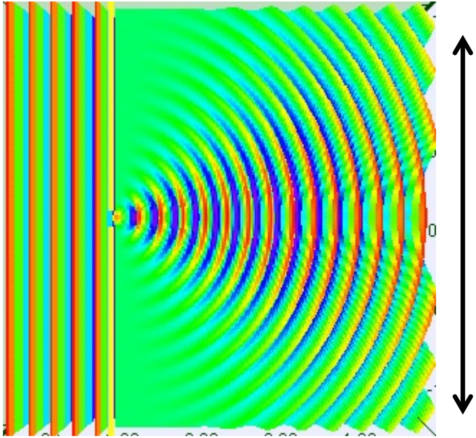

wide diffraction pattern

small pinhole

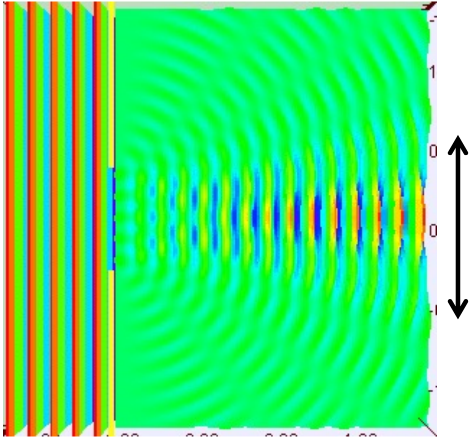

narrow diffraction pattern

large pinhole

What about light efficiency?

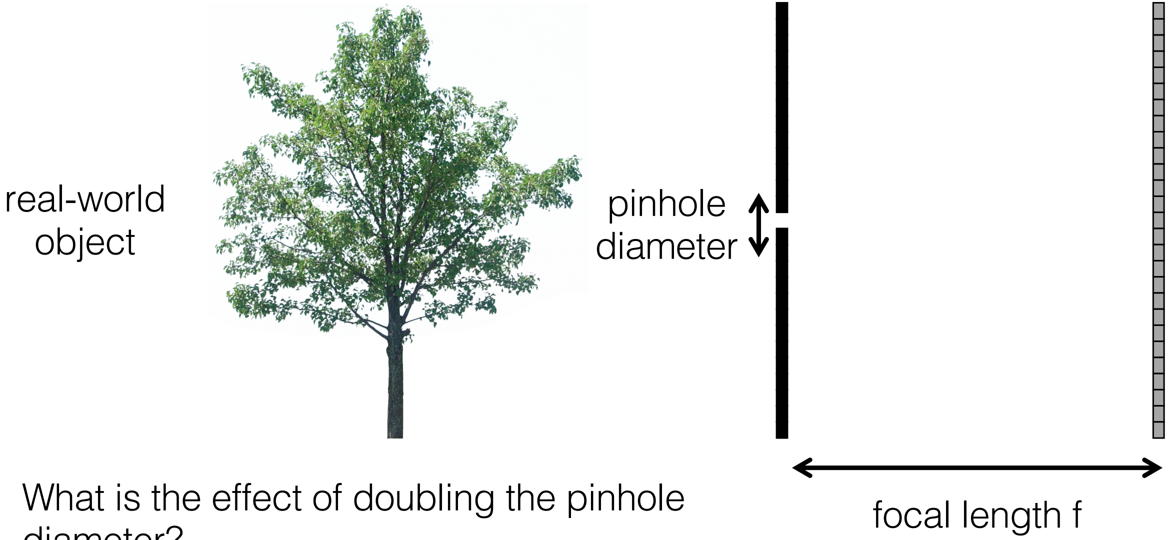

- What is the effect of doubling the pinhole diameter?
- What is the effect of doubling the focal length?

What about light efficiency?

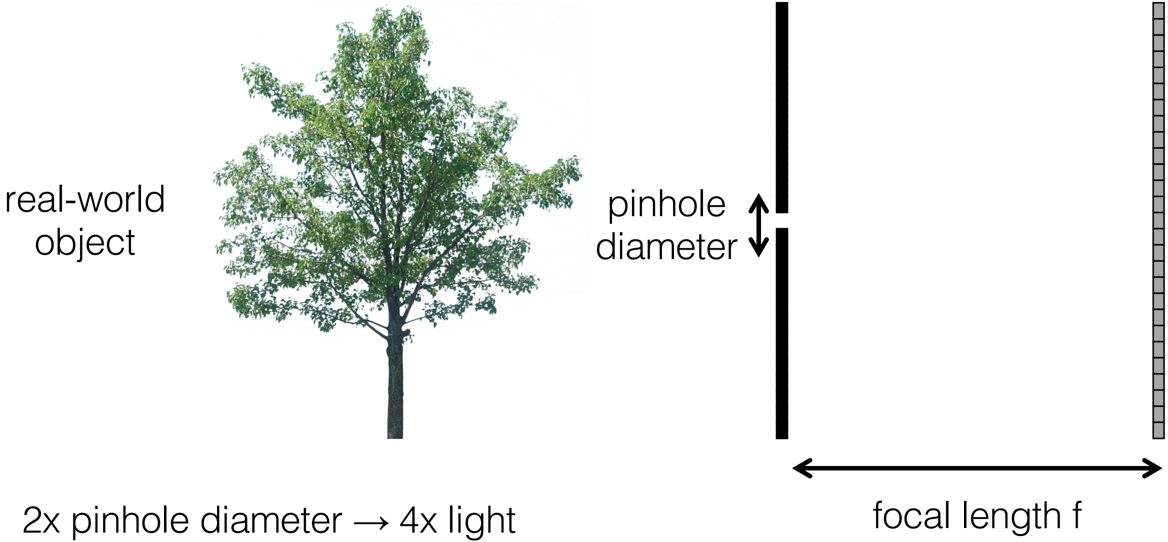

- 2x pinhole diameter → 4x light
- 2x focal length → ¼x light

Pinhole Camera / Camera Obscura

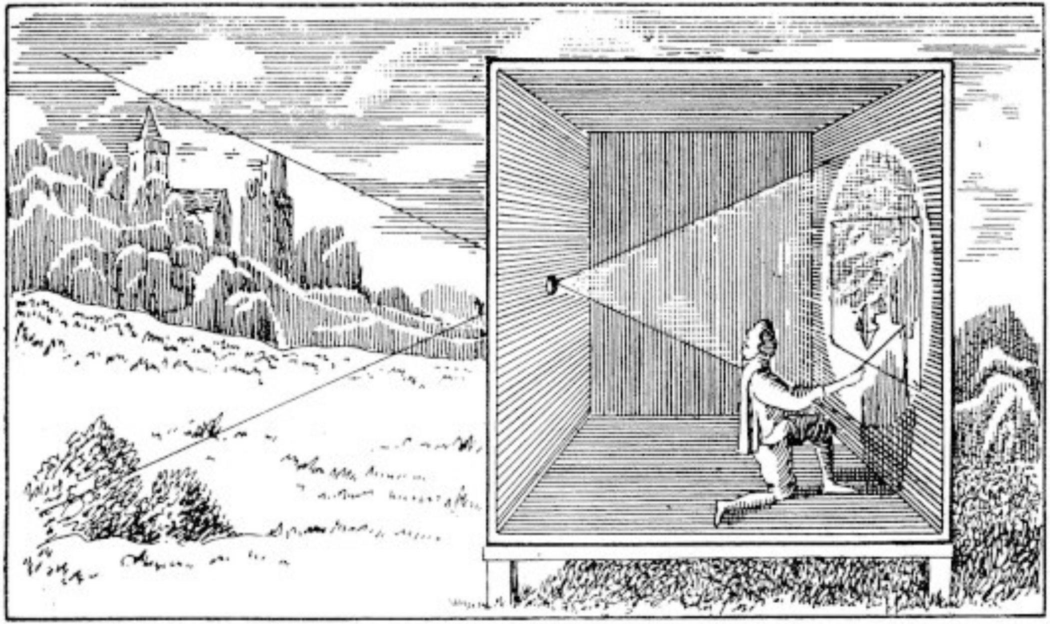

Mo-Ti (Chinese Philosopher) 470-390 BC

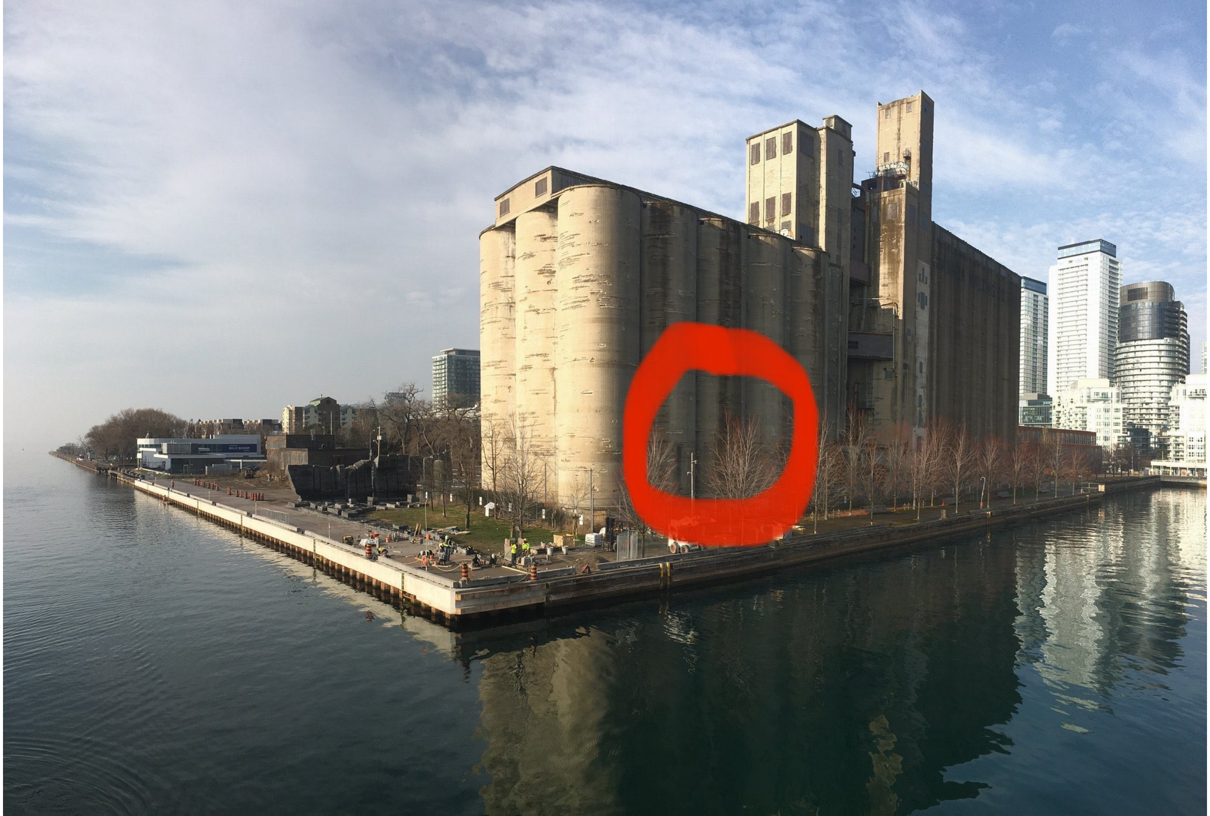

Bryan Bowen
@bryanmbowen

Fun discovery - a small crack in the eastern facade of the Canada Malting Co silos has created a perfect pinhole camera. The result: real time projection of Toronto’s waterfront on the silo’s interior curved surfaces. An unplugged projection show!

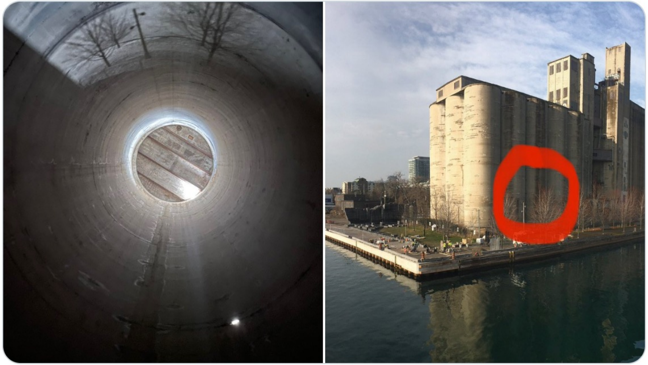

9:37 AM · Jan 27, 2022 · Twitter for iPhone

656 Retweets **70** Quote Tweets **2,836** Likes

Abelardo Morell

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

SLR Camera

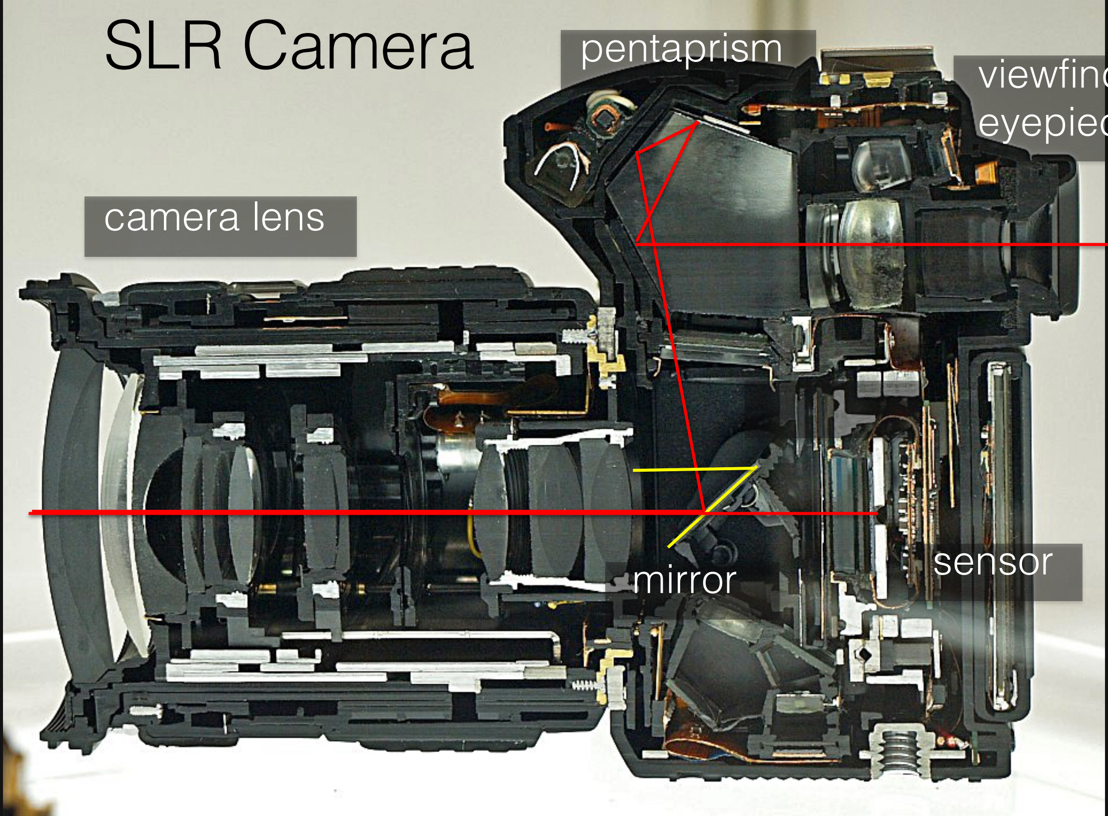
pentaprism
viewfinder
eyepiece
camera lens
mirror
sensor

SLR Camera

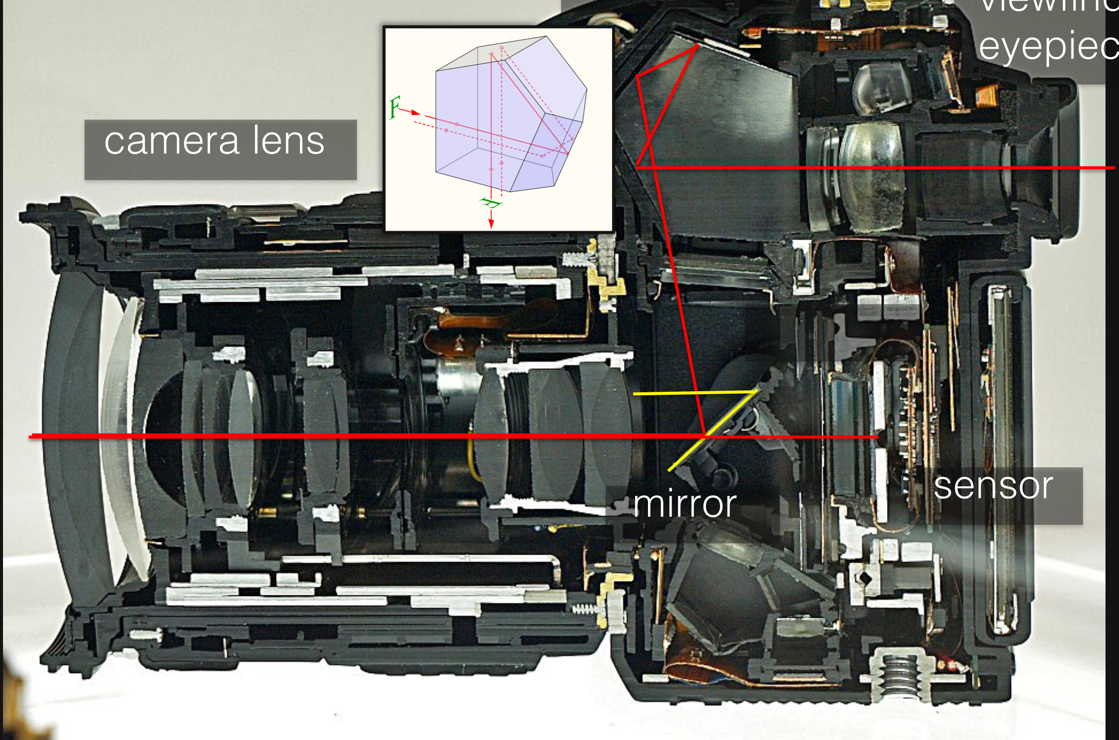
pentaprism
viewfinder eyepiece
camera lens
mirror
sensor

Niepce “View from the Window at Le Gras”, 1826

1826
8h exp

Daguerrotype

- invented in 1836 by Louis Daguerre
- lenses focus light, better chemicals!

Daguerre “Boulevard du Temple”, 1838

exposure
10-12 mins

Lenses

- focus light
- magnify objects

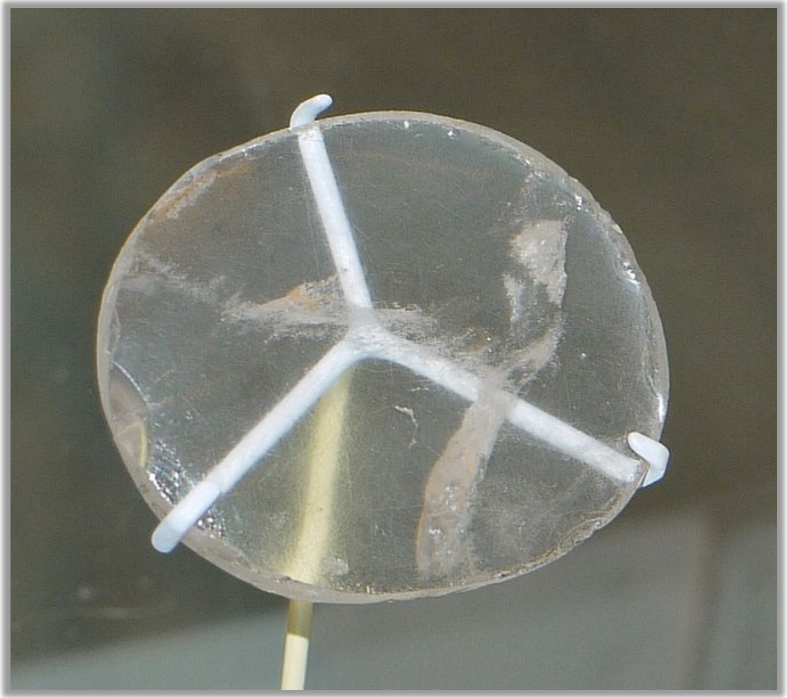

Nimrud lens - 2700 years old

What is a lens?

A piece of glass manufactured to have a specific shape

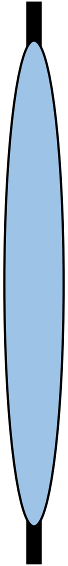

What is a lens?

A piece of glass manufactured to have a specific shape

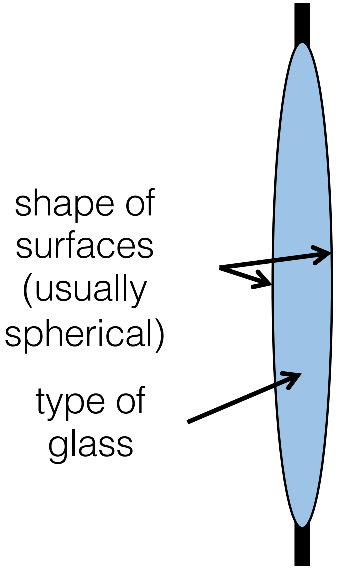

What is a lens?

A piece of glass manufactured to have a specific shape

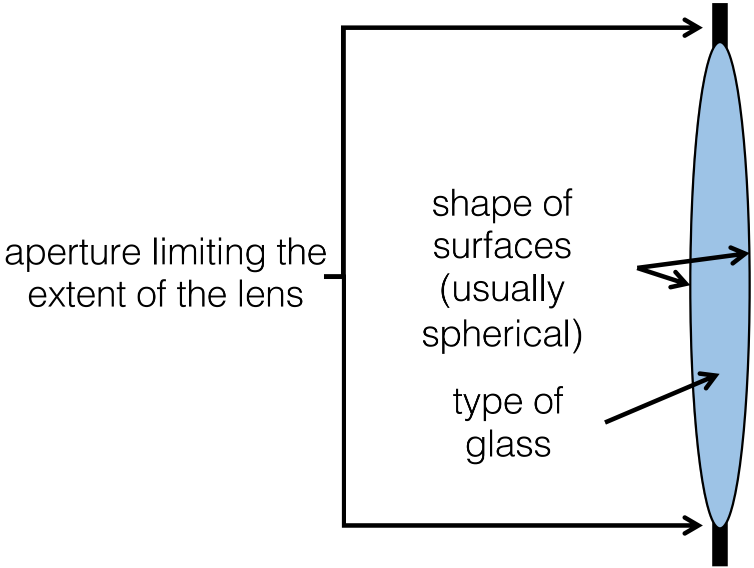

What is a lens?

A piece of glass manufactured to have a specific shape

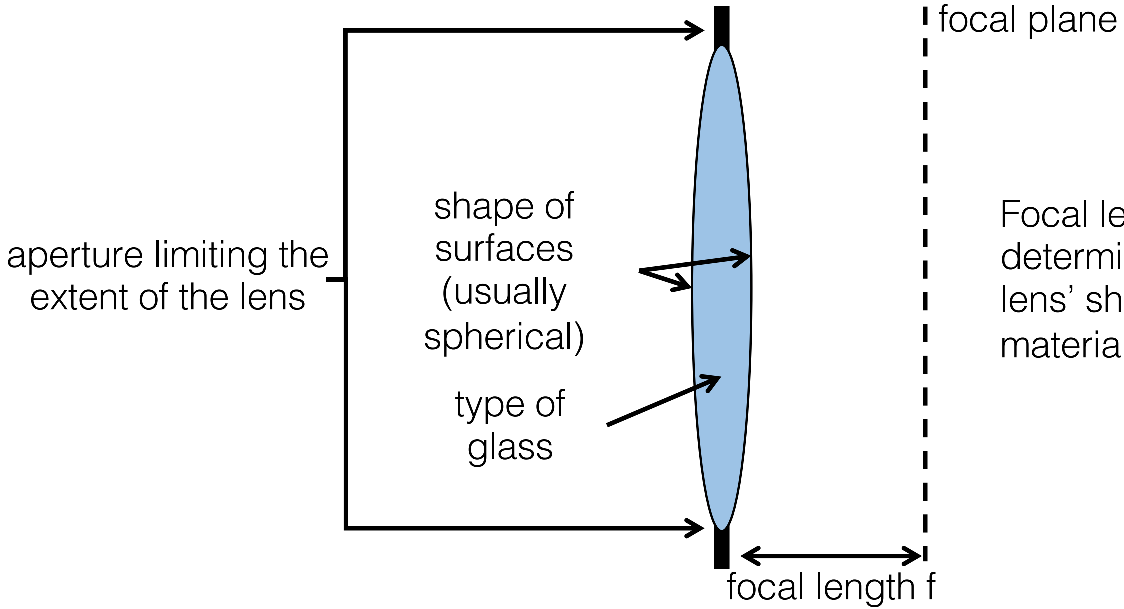

Focal length is determined by the lens' shape and material

How does a lens work?

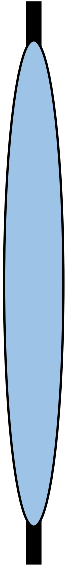

Refraction

Refraction is the bending of rays of light when they move from one material to another

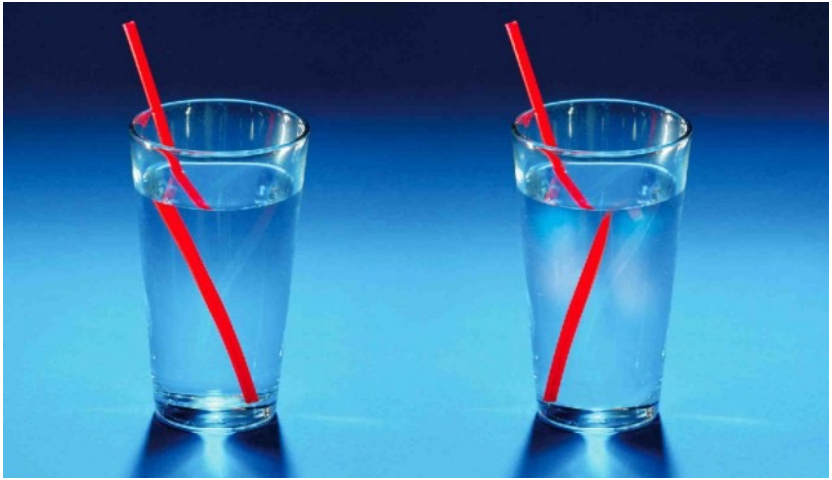

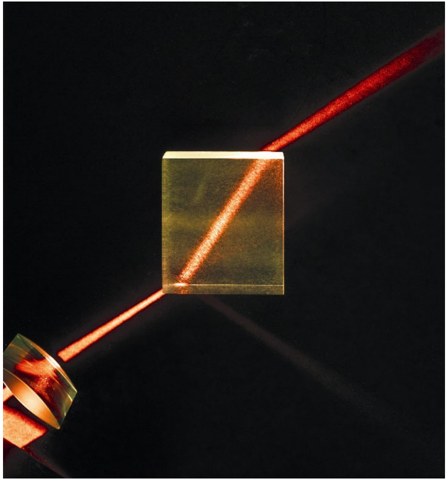

How does a lens work?

Lenses are designed so that their refraction makes light rays bend in a very specific way.

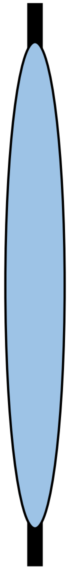

Thin lens model

Simplification of geometric optics for <u>well-designed</u> lenses.

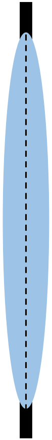

Thin lens model

Simplification of geometric optics for well-designed lenses.

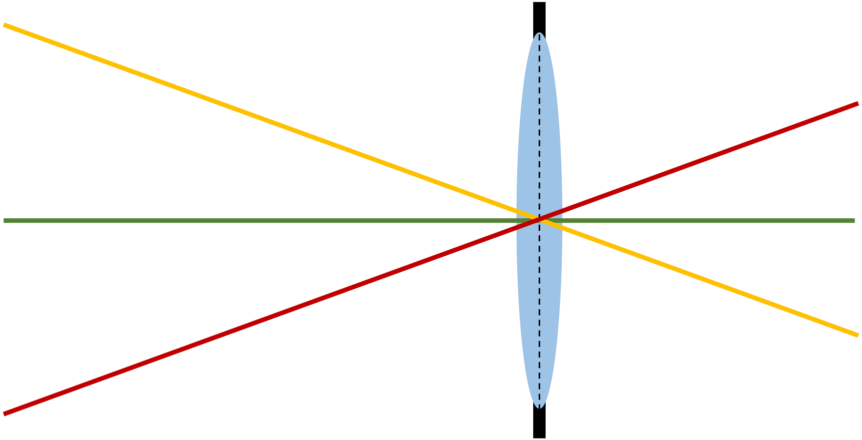

Two assumptions:

1. Rays passing through lens center are unaffected.

Thin lens model

Simplification of geometric optics for well-designed lenses.

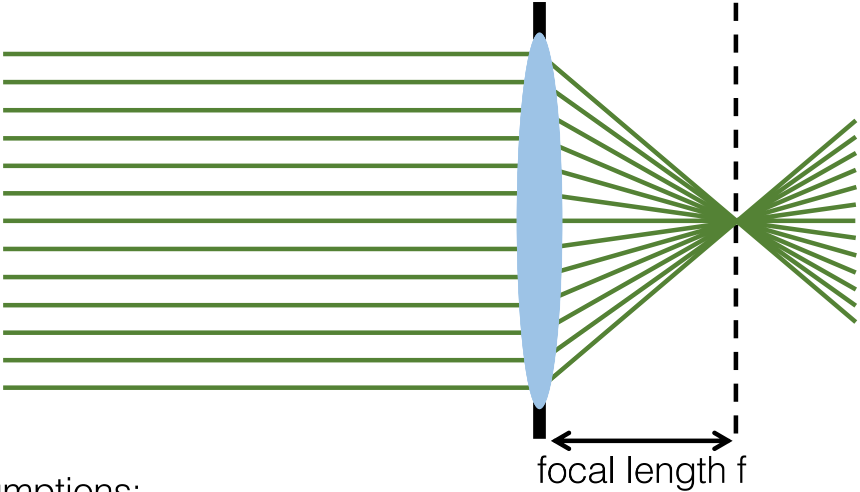

Two assumptions:

1. Rays passing through lens center are unaffected.
2. Parallel rays converge to a single point located on focal plane.

Thin lens model

Simplification of geometric optics for well-designed lenses.

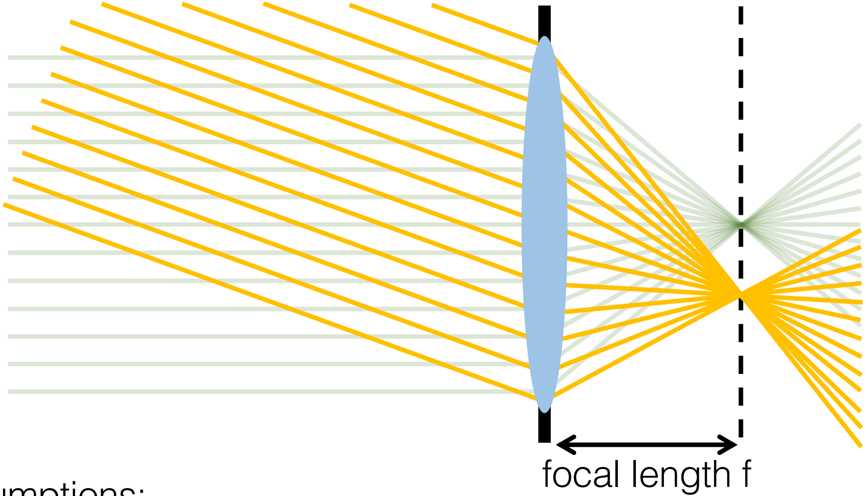

Two assumptions:

1. Rays passing through lens center are unaffected.
2. Parallel rays converge to a single point located on focal plane.

Thin lens model

Simplification of geometric optics for well-designed lenses.

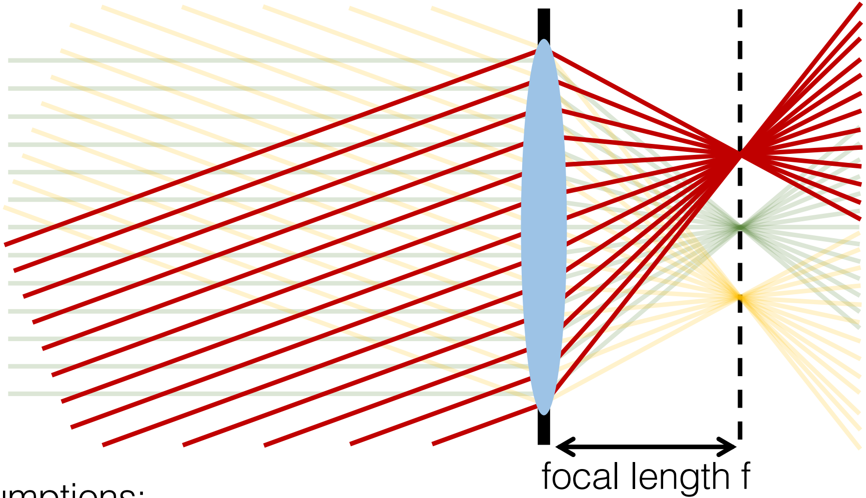

Two assumptions:

1. Rays passing through lens center are unaffected.
2. Parallel rays converge to a single point located on focal plane.

Thin Lens Model

Ray tracing example

- Parallel rays map to the focal plane

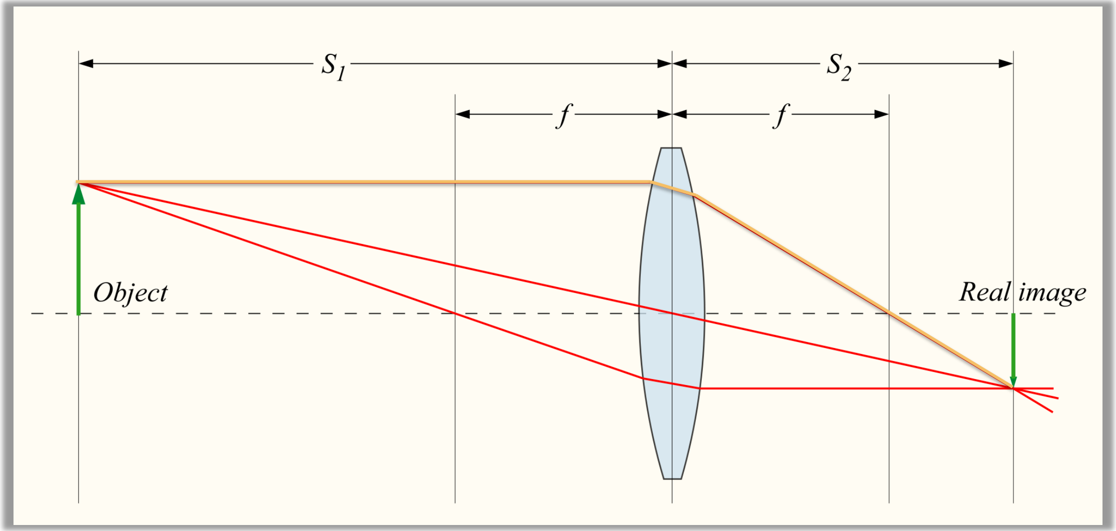

Thin Lens Model

Ray tracing example

- Parallel rays map to the focal plane
- The chief ray passes straight through the center

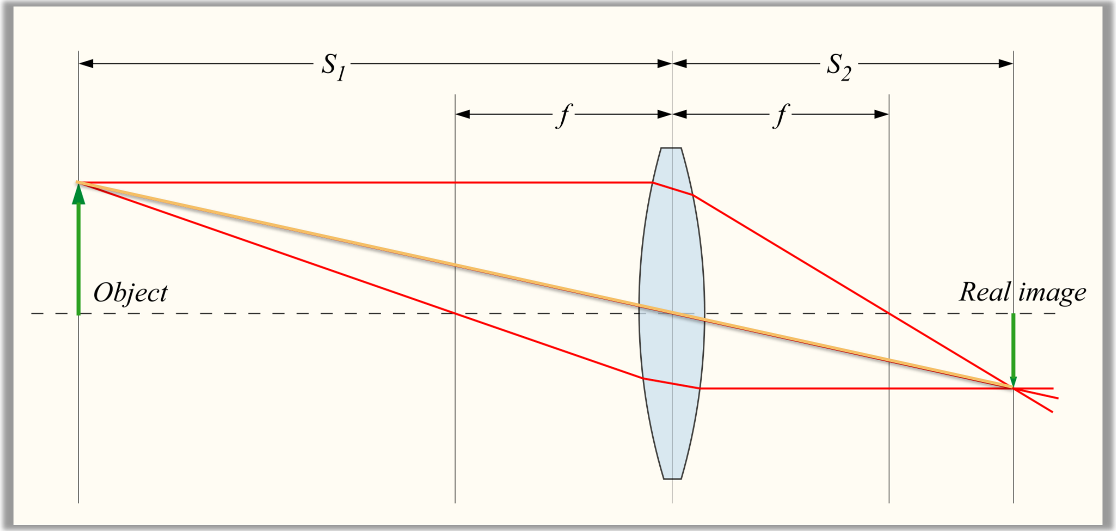

Thin Lens Model

Ray tracing example

- Parallel rays map to the focal plane
- The chief ray passes straight through the center
- The ray that passes through the near focal plane becomes parallel

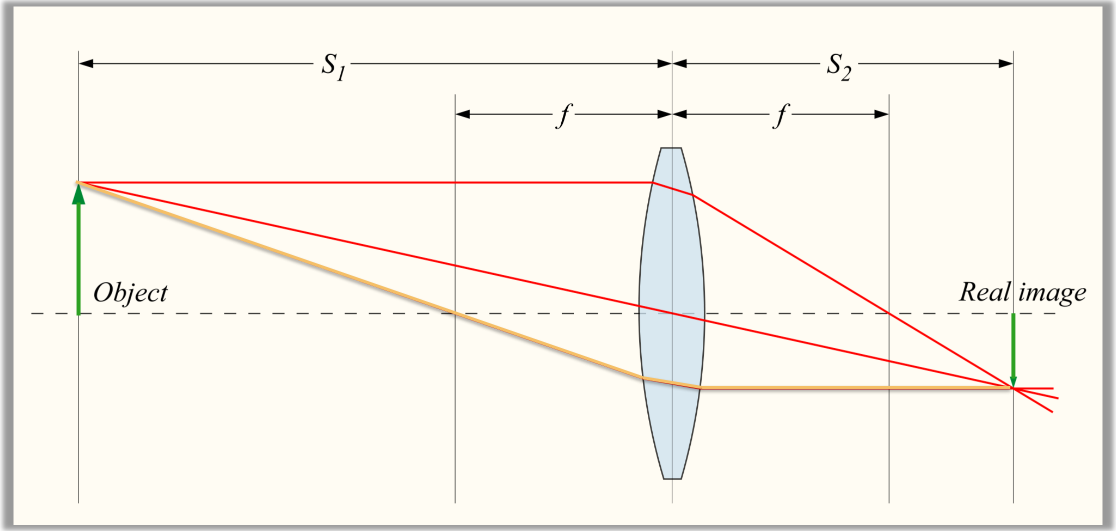

Thin Lens Model

Thin lens formula $$\frac{1}{f} = \frac{1}{S_1} + \frac{1}{S_2}$$

magnification: $$M = -\frac{S_2}{S_1} = \frac{f}{f - S_1}$$

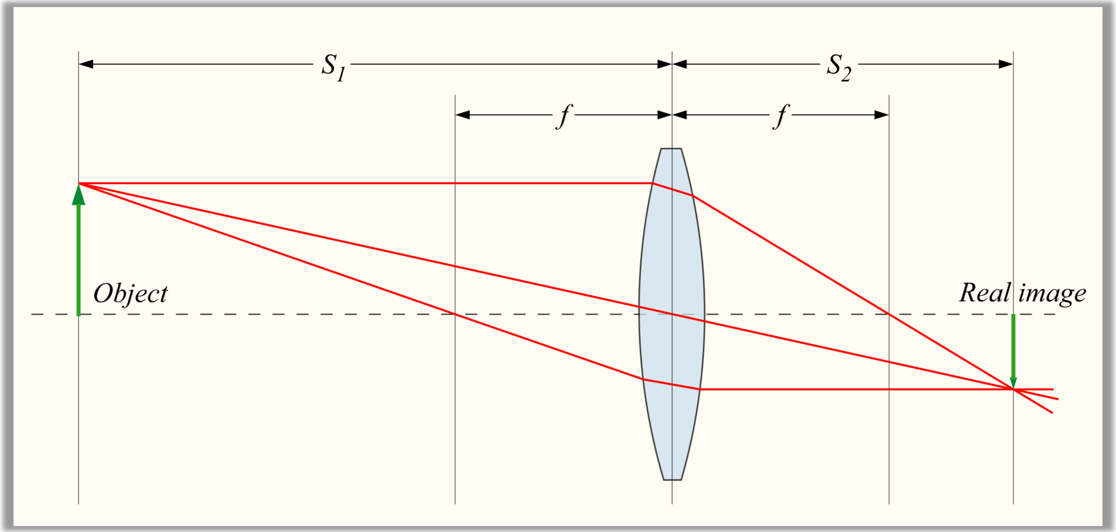

Convex Lens

- f > 0

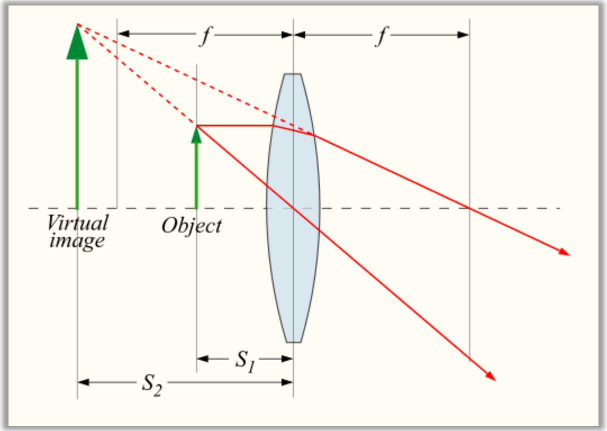

Convex Lens

- f > 0
- S1<f: magnifying glass

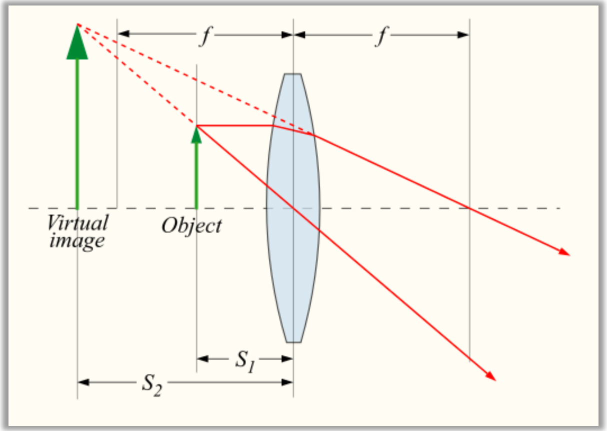

Convex Lens

- f > 0
- S1<f: magnifying glass
- S2 negative for virtual image
- S1 positive for object/real image

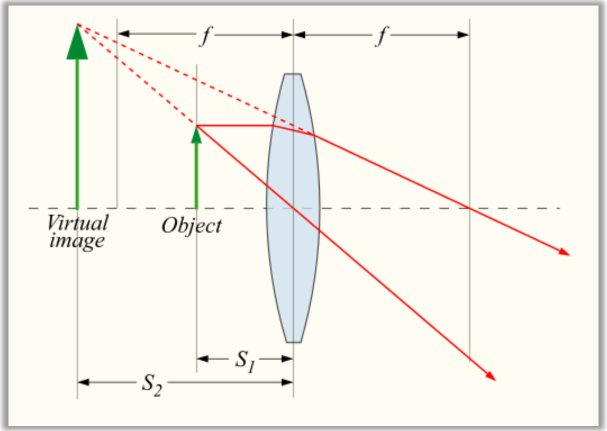

Convex Lens

- f > 0
- S1<f: magnifying glass
- S2 negative for virtual image
- S1 positive for object/real image

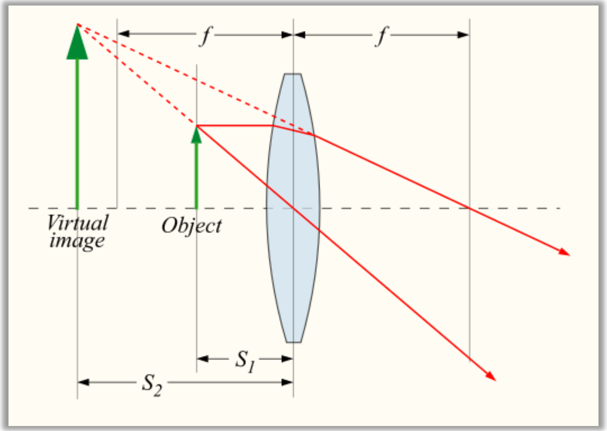

moves object further (far-sighted)

Convex Lens

- f > 0
- S1<f: magnifying glass
- S2 negative for virtual image
- S1 positive for object/real image

Concave Lens

- f < 0

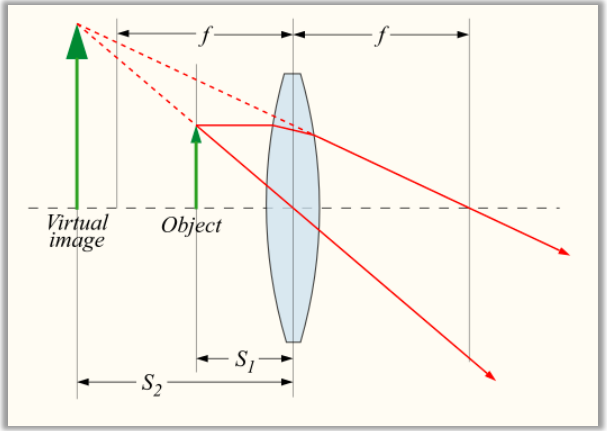

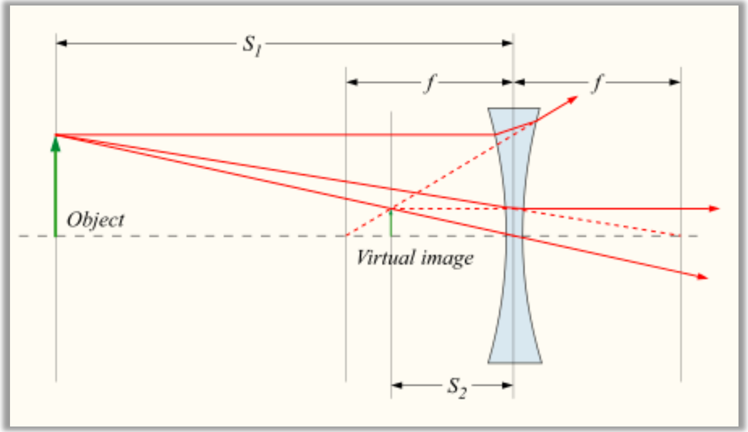

moves object further (far-sighted)

Convex Lens

- f > 0
- S1<f: magnifying glass
- S2 negative for virtual image
- S1 positive for object/real image

Concave Lens

- f < 0
- S1>f: minification

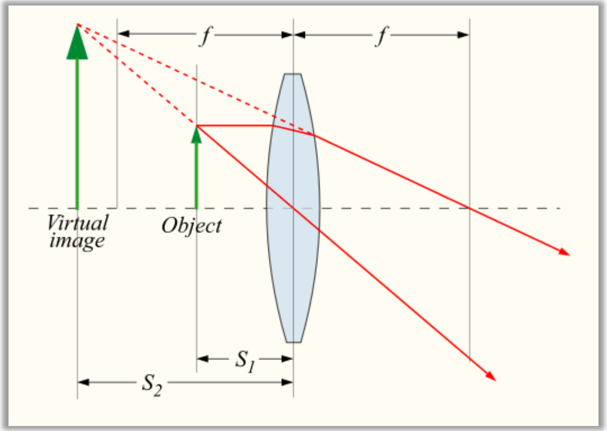

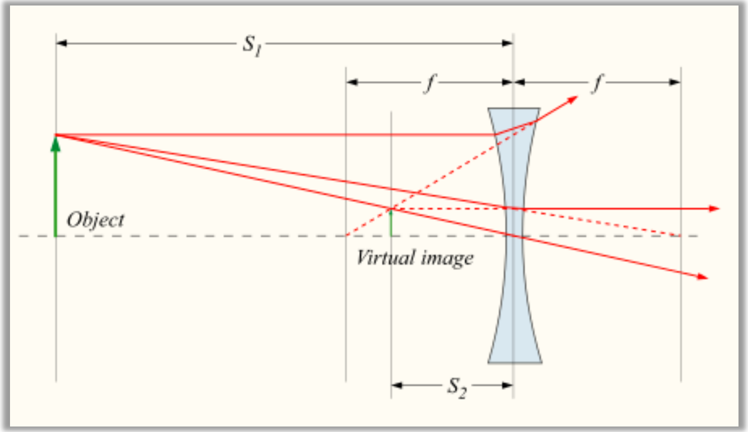

moves object further (far-sighted)

Convex Lens

- f > 0
- S1<f: magnifying glass
- S2 negative for virtual image
- S1 positive for object/real image

Concave Lens

- f < 0
- S1>f: minification
- S2 negative for virtual image
- S1 positive for object/real image

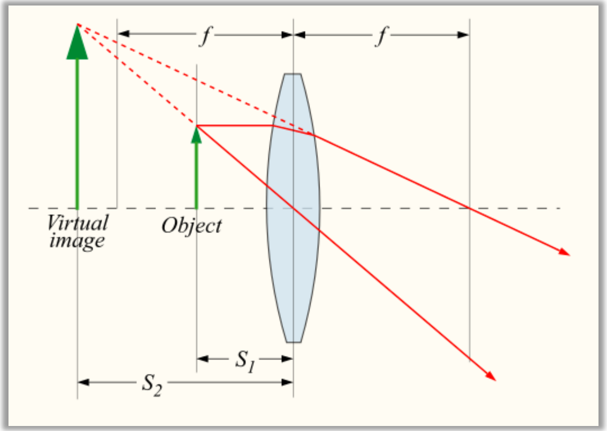

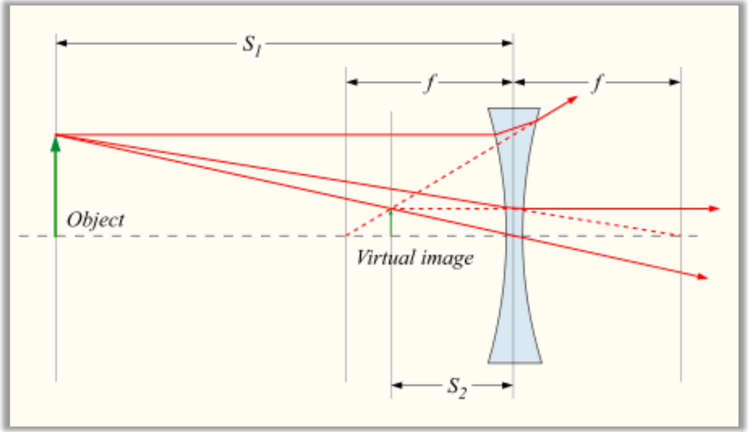

moves object further (far-sighted)

Convex Lens

- $f > 0$
- $S1<f$: magnifying glass
- S2 negative for virtual image
- S1 positive for object/real image

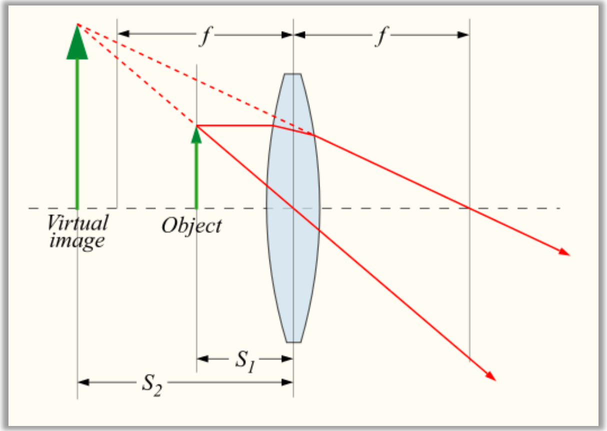

moves object further (far-sighted)

Concave Lens

- $f < 0$
- $S1>f$: minification
- S2 negative for virtual image
- S1 positive for object/real image

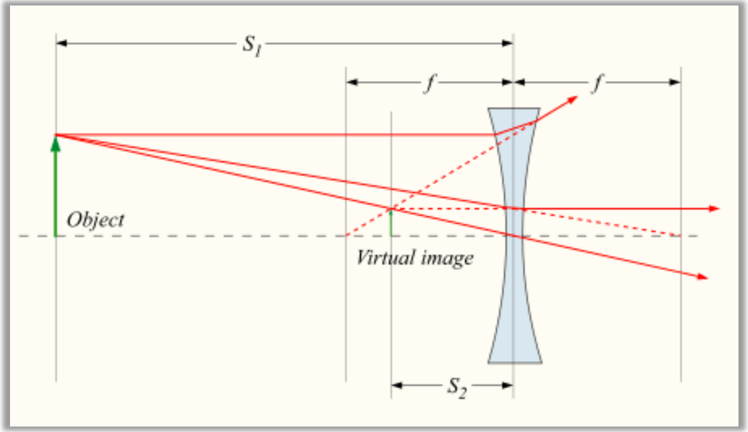

brings object closer (near sighted)

Yes, but…

Thin lenses are a fiction

The thin lens model assumes that the lens has no thickness, but this is rarely true…

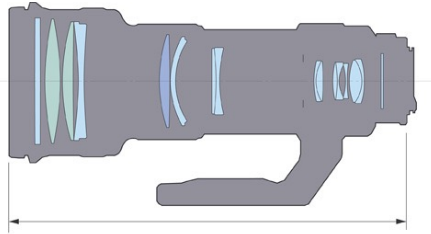

To make real lenses behave like ideal thin lenses, we have to use combinations of multiple lens elements (compound lenses).

Thin lenses are a fiction

The thin lens model assumes that the lens has no thickness, but this is rarely true…

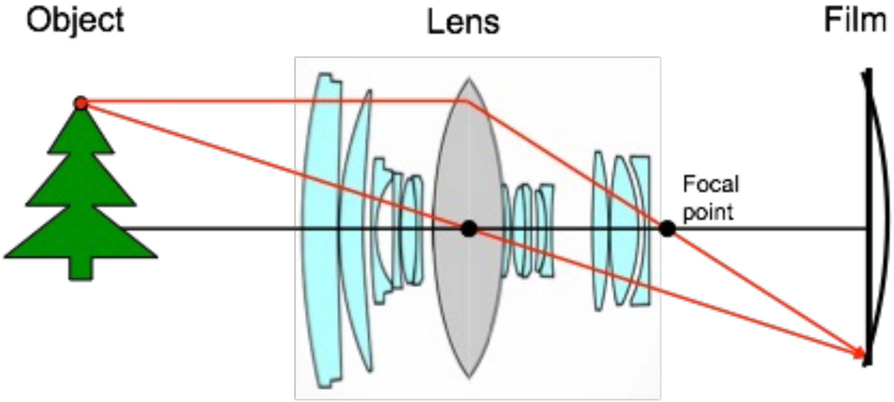

- Even though we have multiple lenses, the entire optical system can be (paraxially) described using a single thin lens of some equivalent focal length and aperture number.

To make real lenses behave like ideal thin lenses, we have to use combinations of multiple lens elements (compound lenses).

Lenses - Aberrations

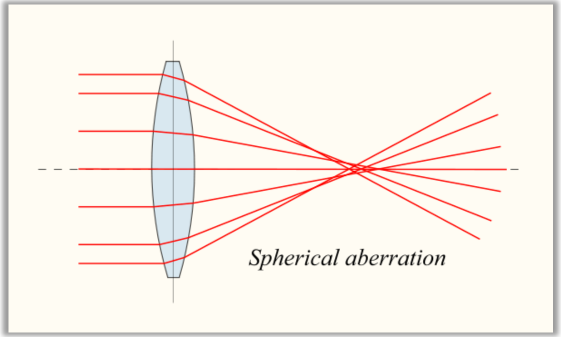

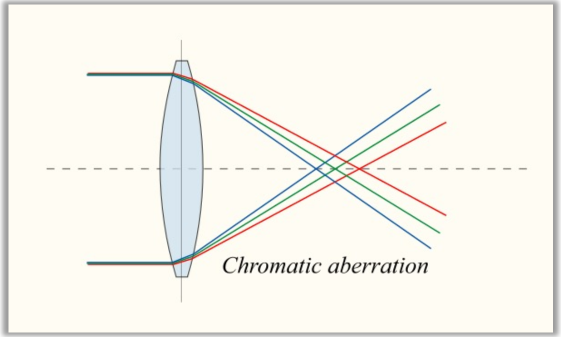

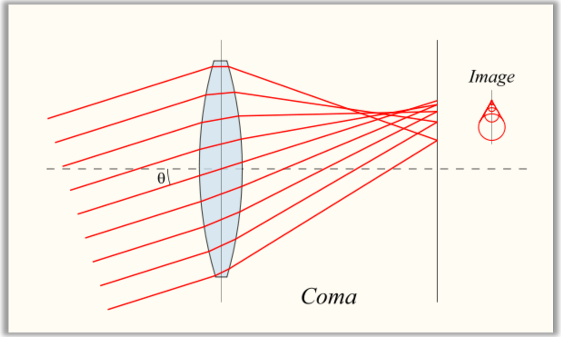

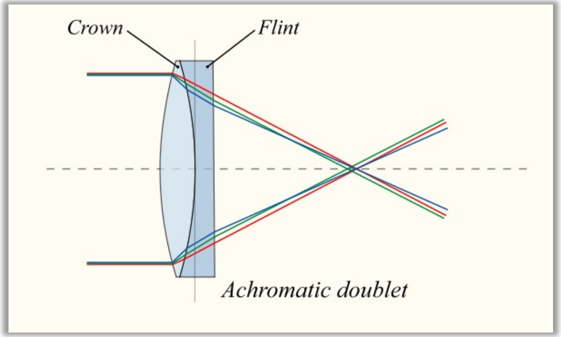

Lenses - Aberrations

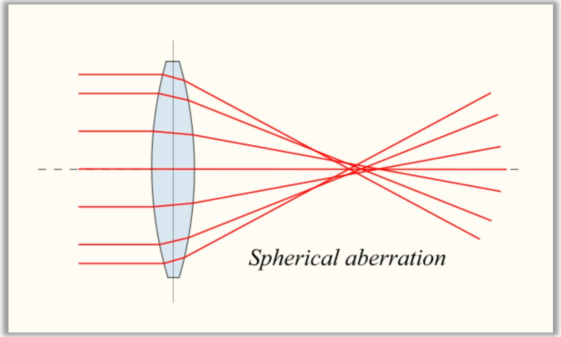

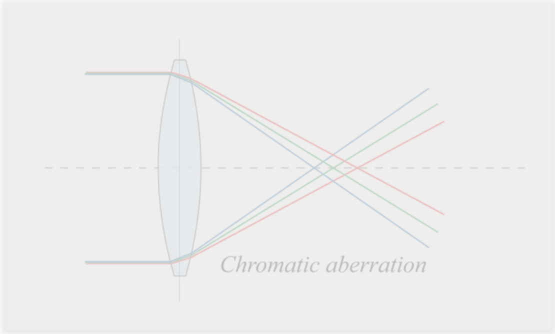

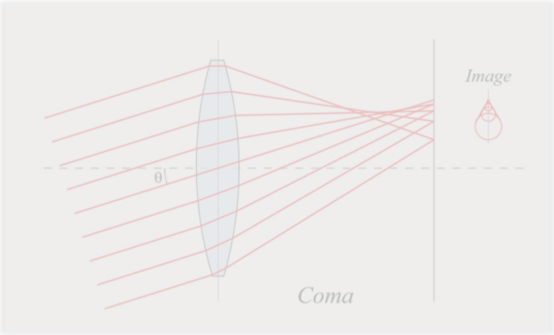

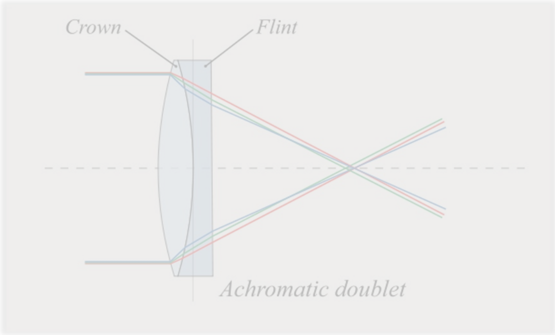

Refraction at interfaces of complicated shapes

What shape should an interface have to make parallel rays converge to a point?

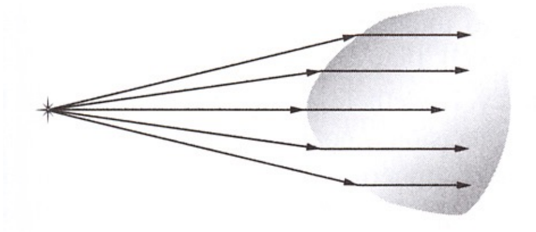

Refraction at interfaces of complicated shapes

What shape should an interface have to make parallel rays converge to a point?

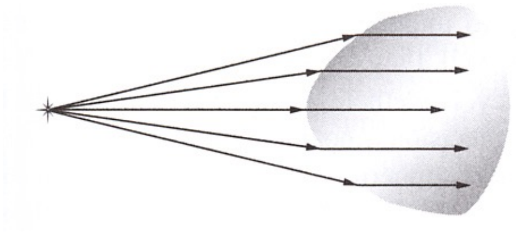

Single hyperbolic interface:
point to parallel rays

Double hyperbolic interface:
point to point rays

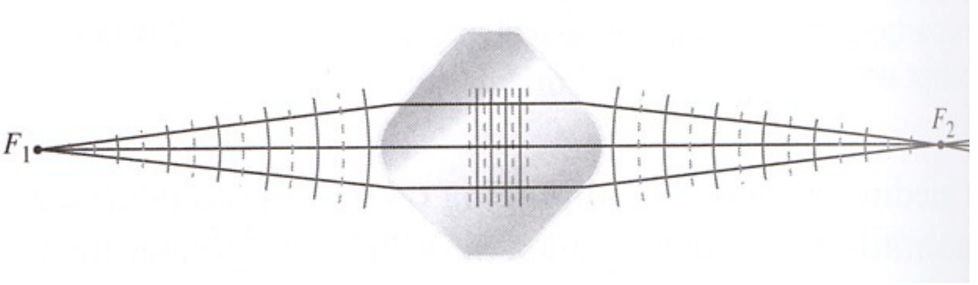

Therefore, lenses should also have hyperbolic shapes.

Spherical lenses

In practice, lenses are often made to have spherical interfaces for ease of fabrication.

- Two roughly fitting curved surfaces ground together will eventually become spherical.

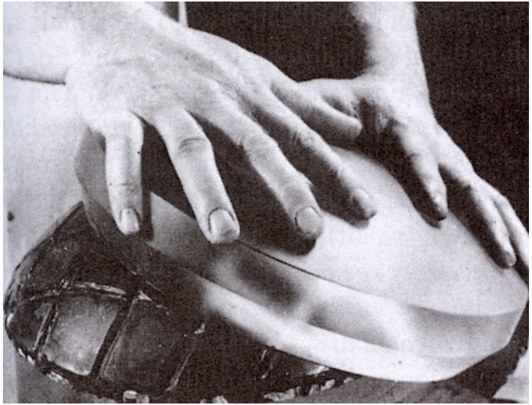

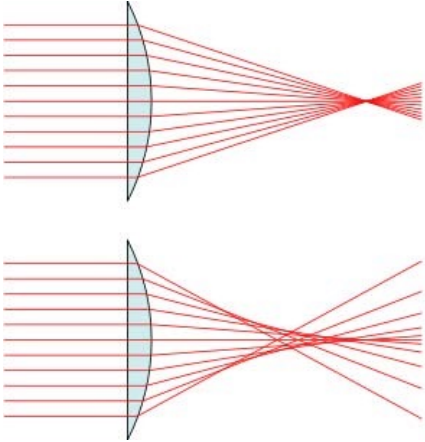

Spherical lenses don't bring parallel rays to a point.

- This is called spherical aberration.
- Approximately axial (i.e., paraxial) rays behave better.

Aberrations

Deviations from ideal thin lens behavior (e.g., imperfect focus).

- Example: spherical aberration.

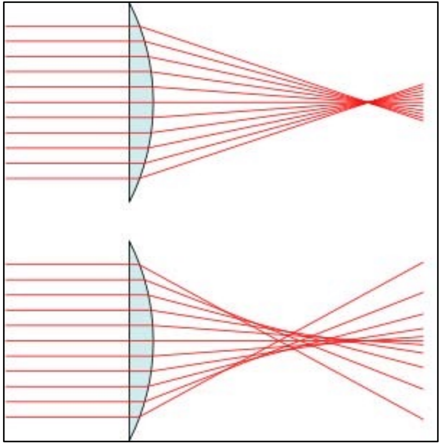

Lenses - Aberrations

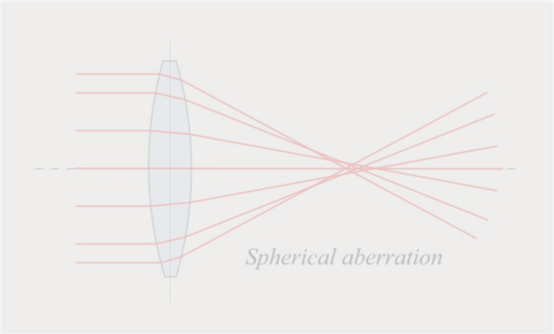

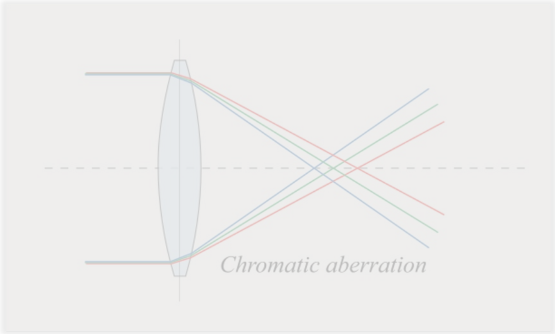

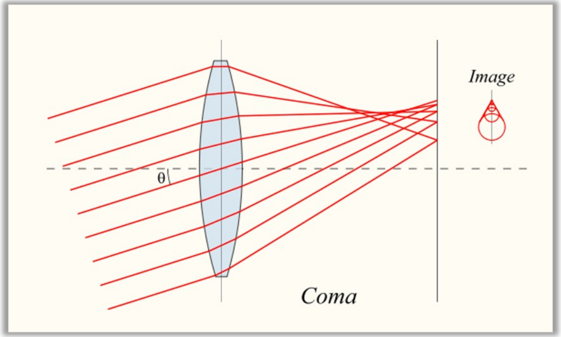

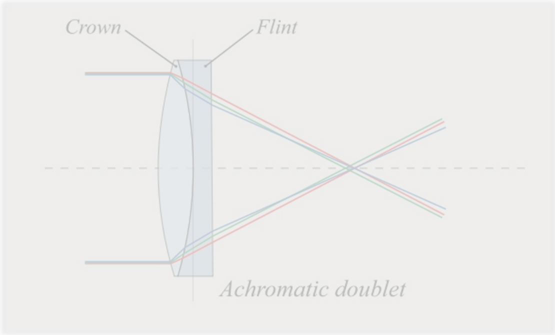

Oblique aberrations

These appear only as we move further from the center of the field of view.

- Contrast with spherical and chromatic, which appear everywhere.
- Many other examples (astigmatism, field curvature, etc.).

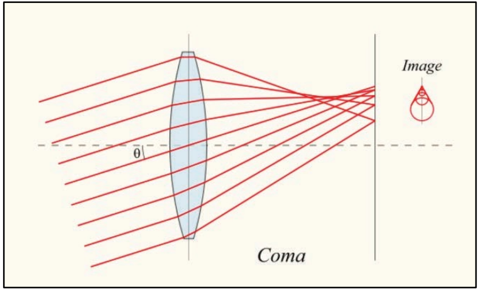

Coma

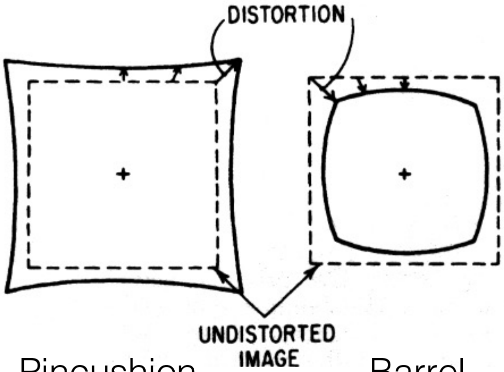

Pincushion Barrel

Distortion example

Lenses - Aberrations

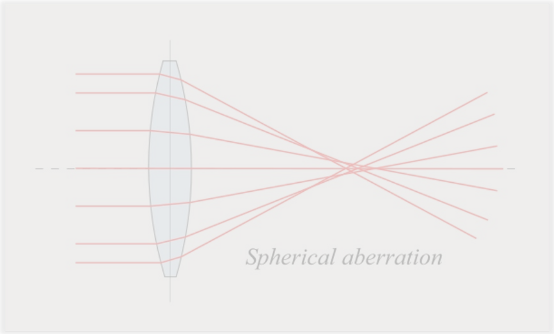

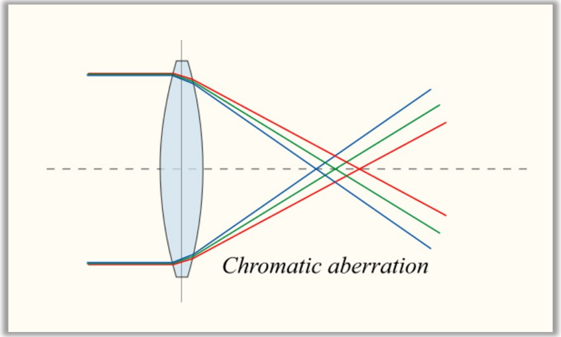

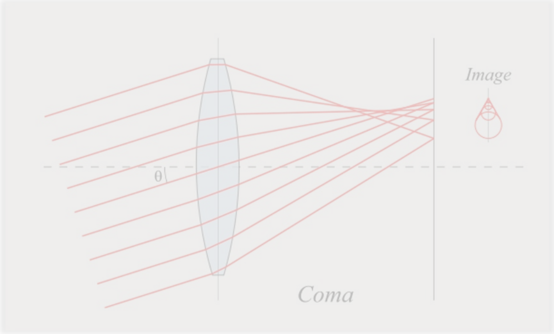

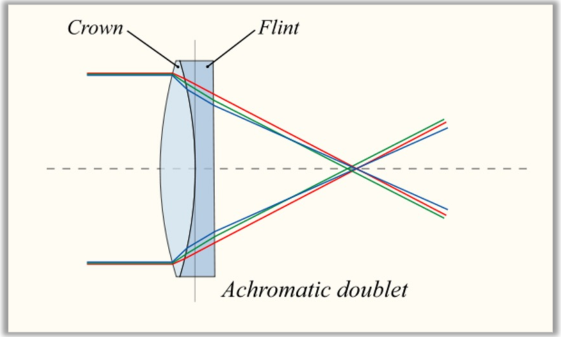

Aberrations

Deviations from ideal thin lens behavior (e.g., imperfect focus).

- Example: chromatic aberration.

focal length shifts with wavelength

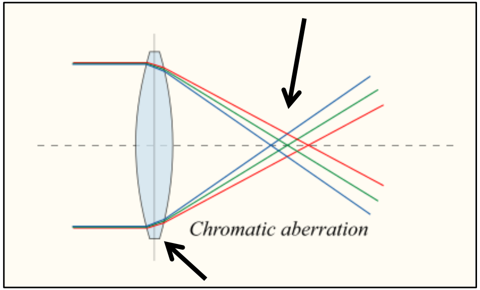

glass has dispersion (refractive index changes with wavelength)

one lens cancels out dispersion of other

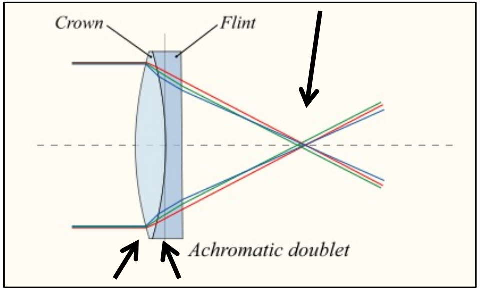

glasses of different refractive index

Using a doublet (two-element compound lens), we can reduce chromatic aberration.

Chromatic aberration examples

A costly aberration

Hubble telescope originally suffered from severe spherical aberration.

- COSTAR mission inserted optics to correct the aberration.

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

Field of View

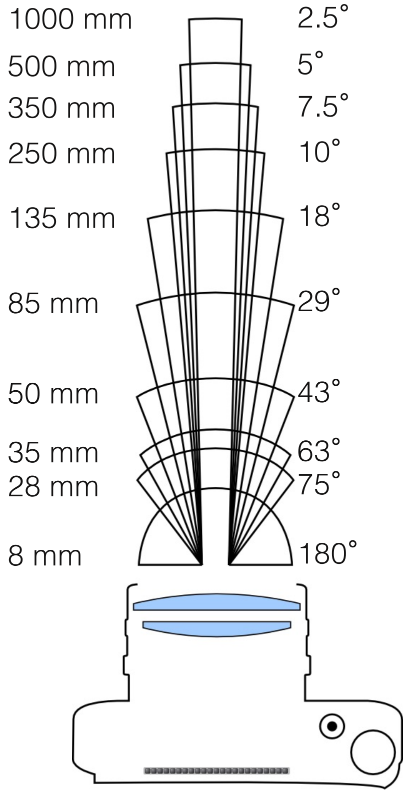

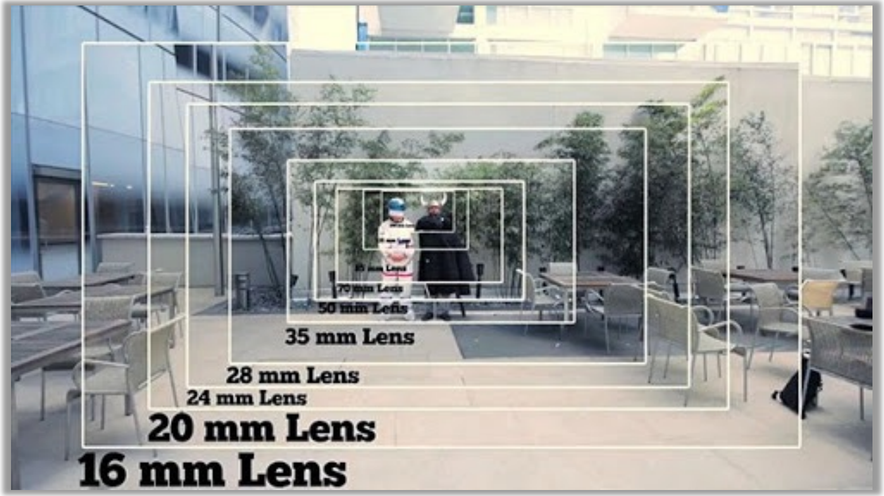

Andrew McWilliams

Field of View

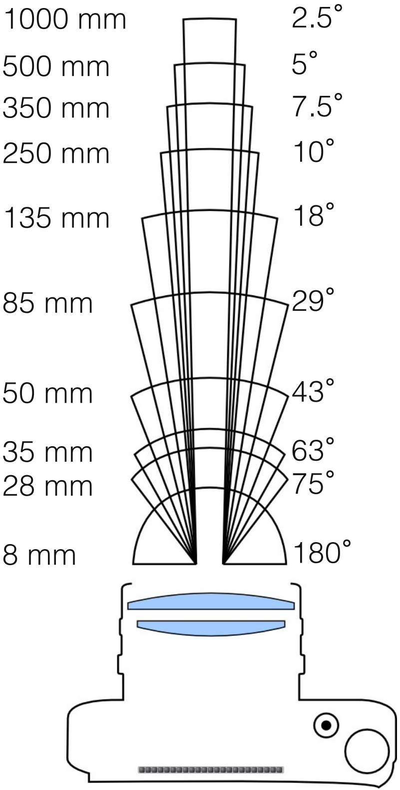

Hubble – what's the focal length?

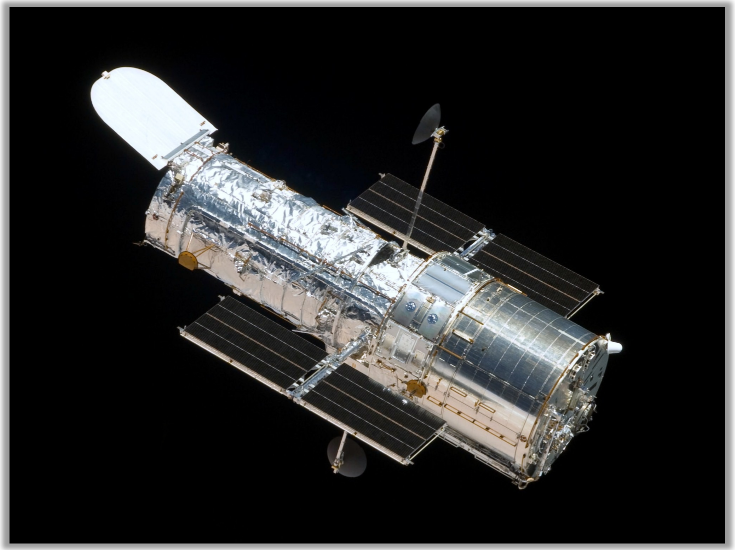

Field of View

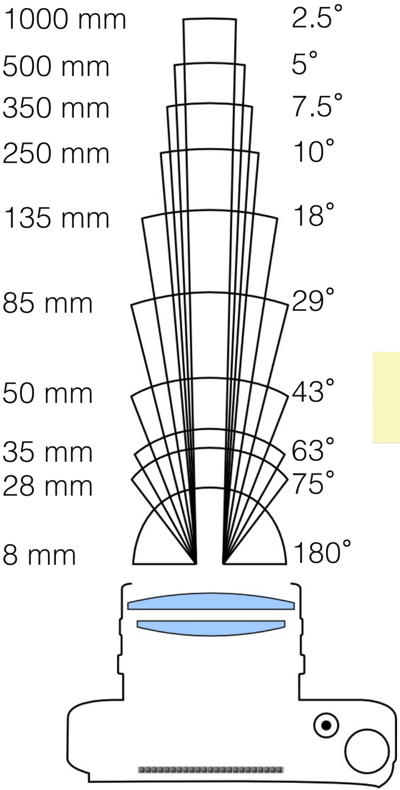

Hubble – what's the focal length?

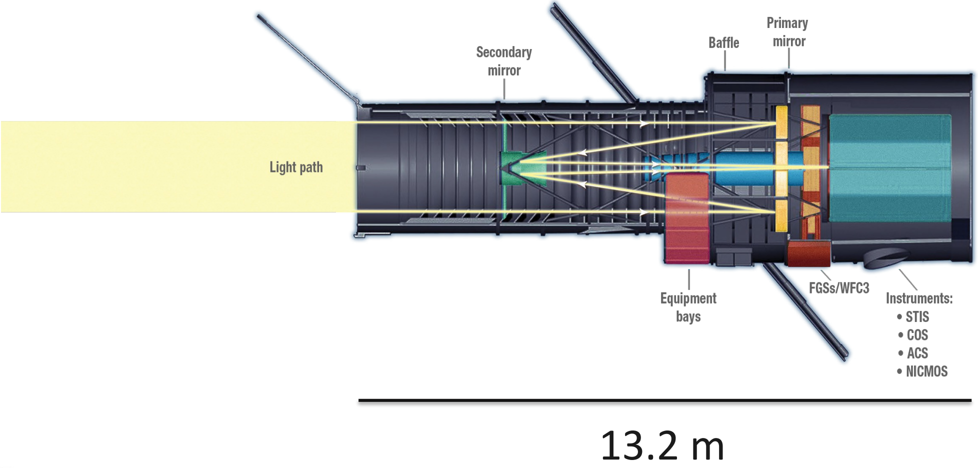

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

Aperture

- aperture controls amount of light

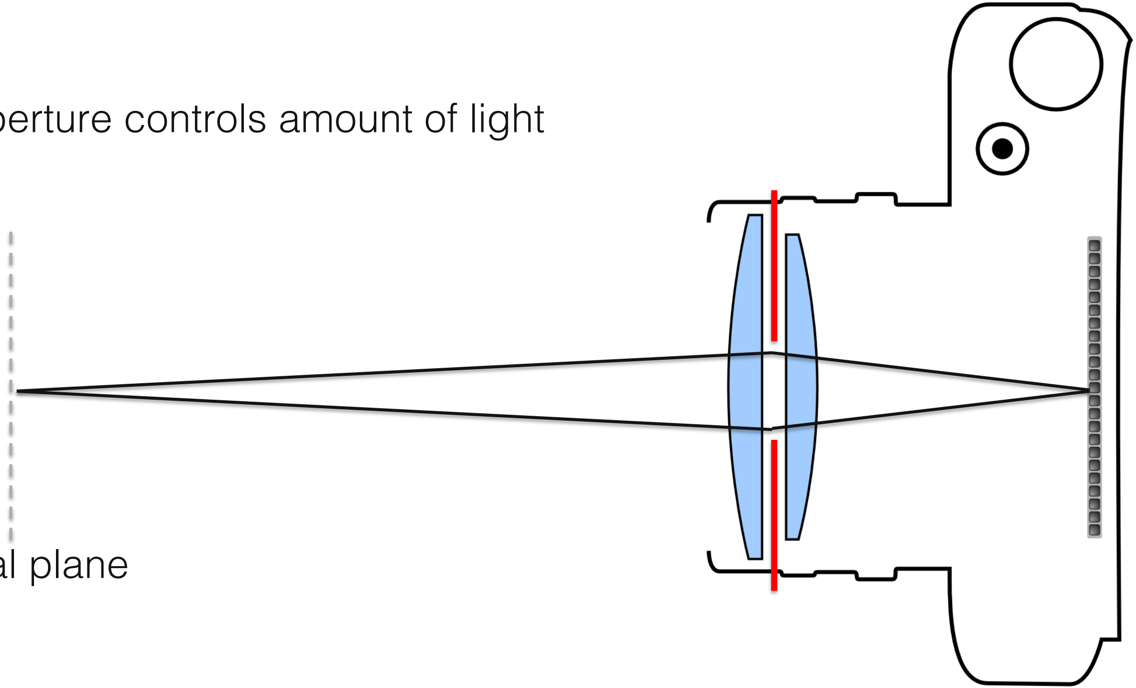

Aperture

- aperture controls amount of light

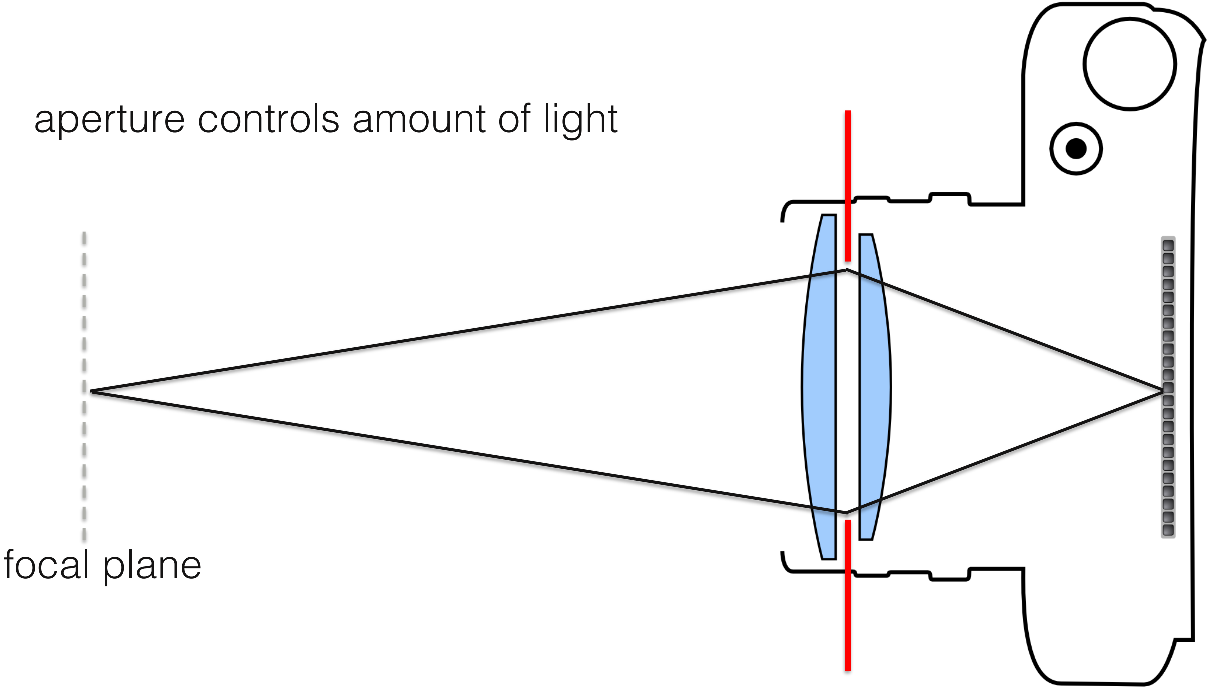

Aperture size

Most lenses have variable aperture size.

- F-number notation: "f/1.4" means $f/=1.4$ (**focal length / diameter**).
- Usually aperture sizes available at steps of one-half or one-third stops.
- Older lenses have separate manual aperture ring.
- Modern lenses control the aperture through a dial on the camera body ("gelded" lenses).

f/1.4

f/2.8

f/4

f/8

f/16

Aperture size

Most lenses have variable aperture size.

- F-number notation: “f/1.4” means $f/ = 1.4$ (**focal length / diameter**).
- Usually aperture sizes available at steps of one-half or one-third stops.
- Older lenses have separate manual aperture ring.
- Modern lenses control the aperture through a dial on the camera body (“gelded” lenses).

f/1.4 f/2.8 f/4 f/8 f/16

Reminder: A “stop” changes the amount of light by a factor of 2.

Aperture

- out of focus blur

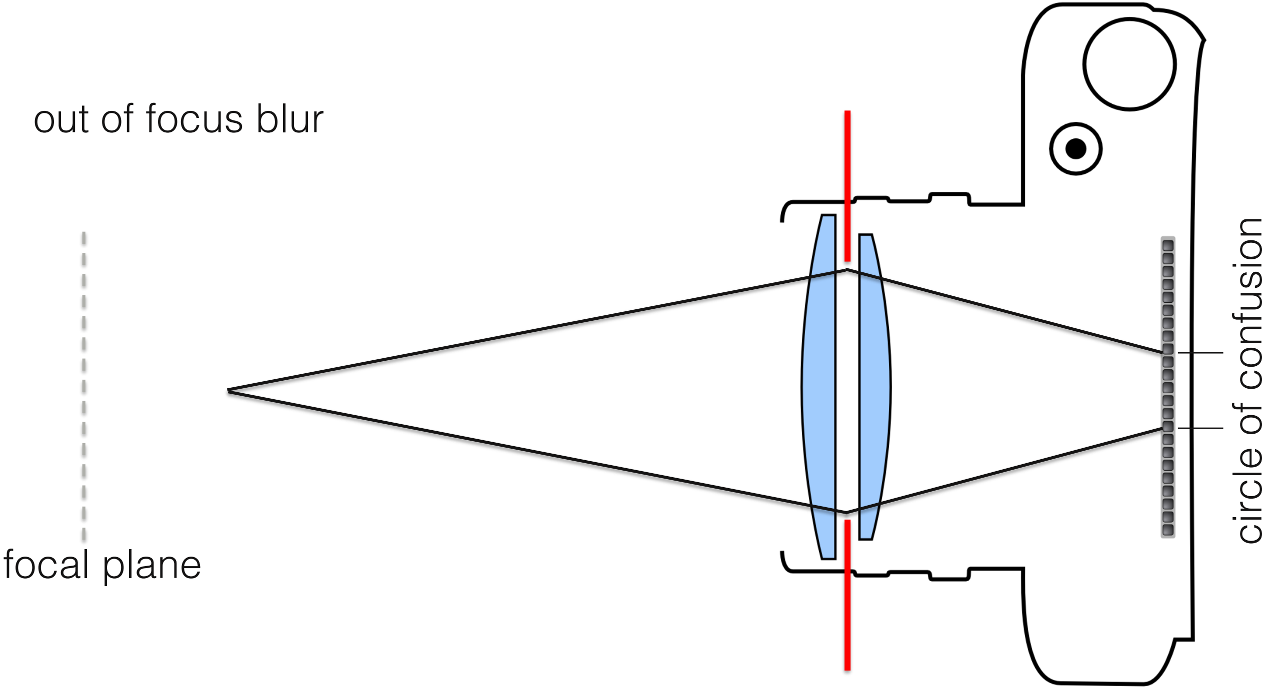

Aperture

- out of focus blur

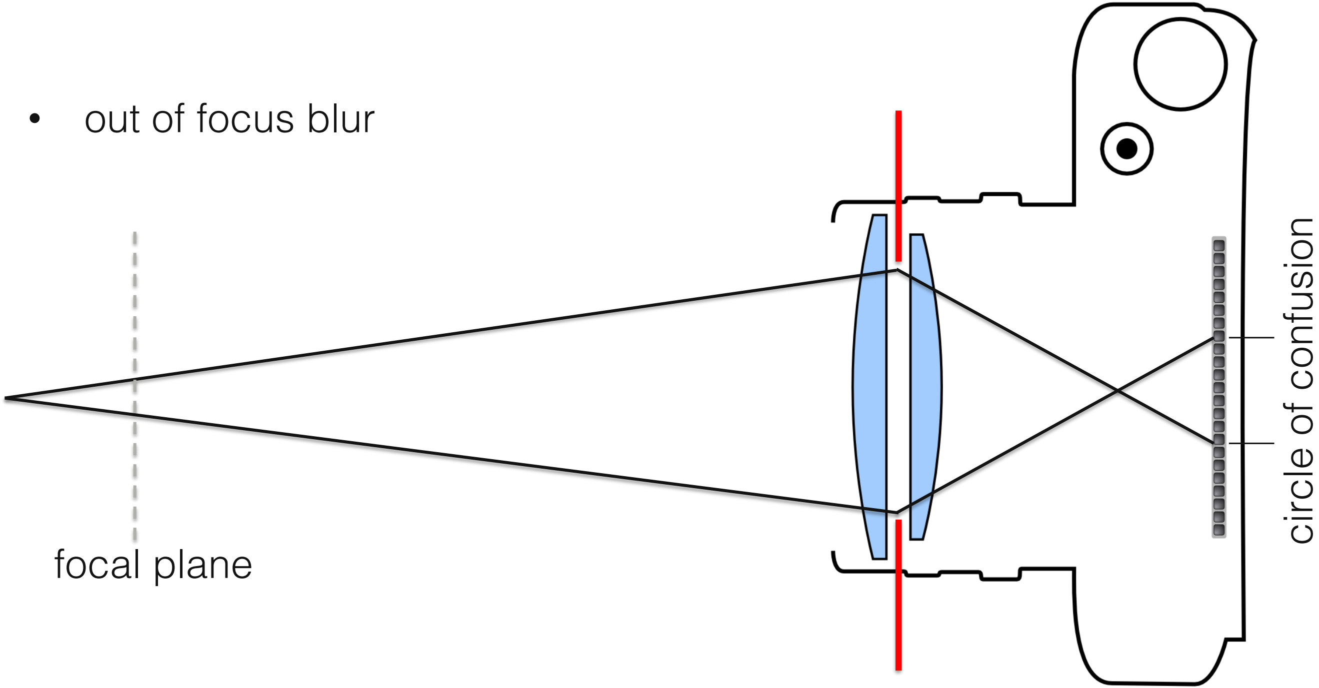

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

Depth of Field

depth of field: max circle of confusion

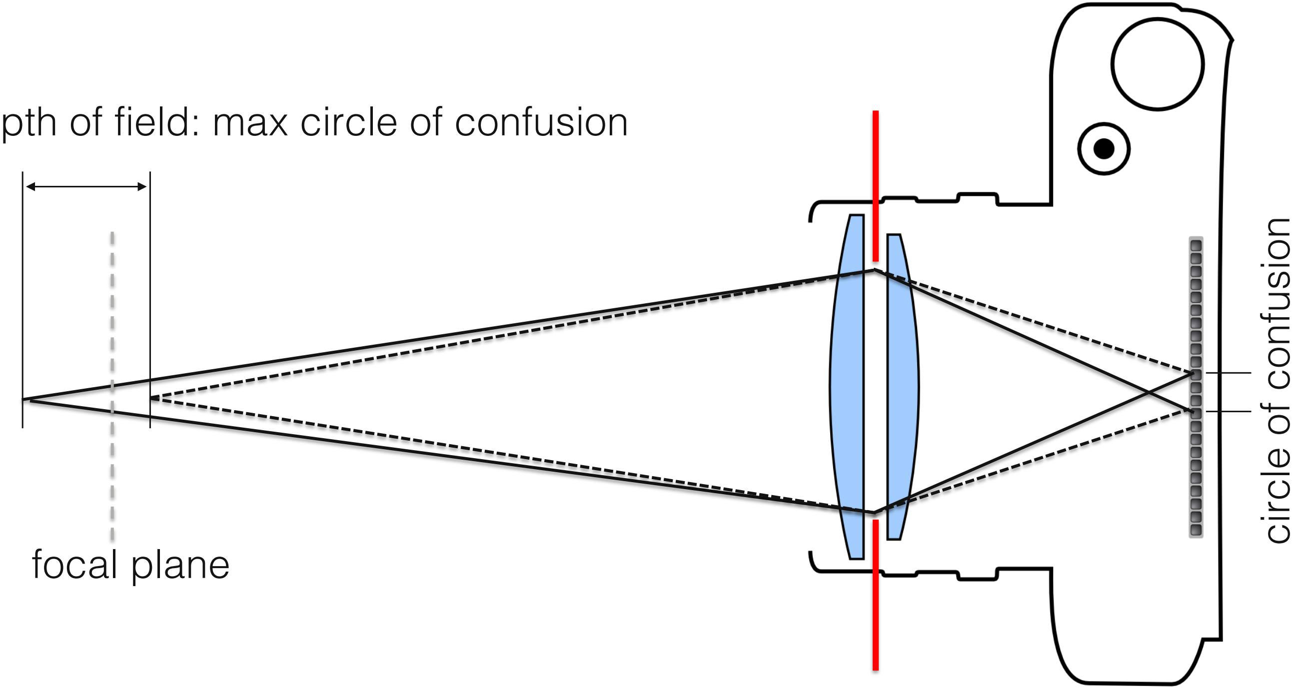

focal plane

Circle of Confusion

$$c = M \cdot D \cdot \frac{|S - S_1|}{S}$$

$$M = \frac{f}{S_1 - f}$$

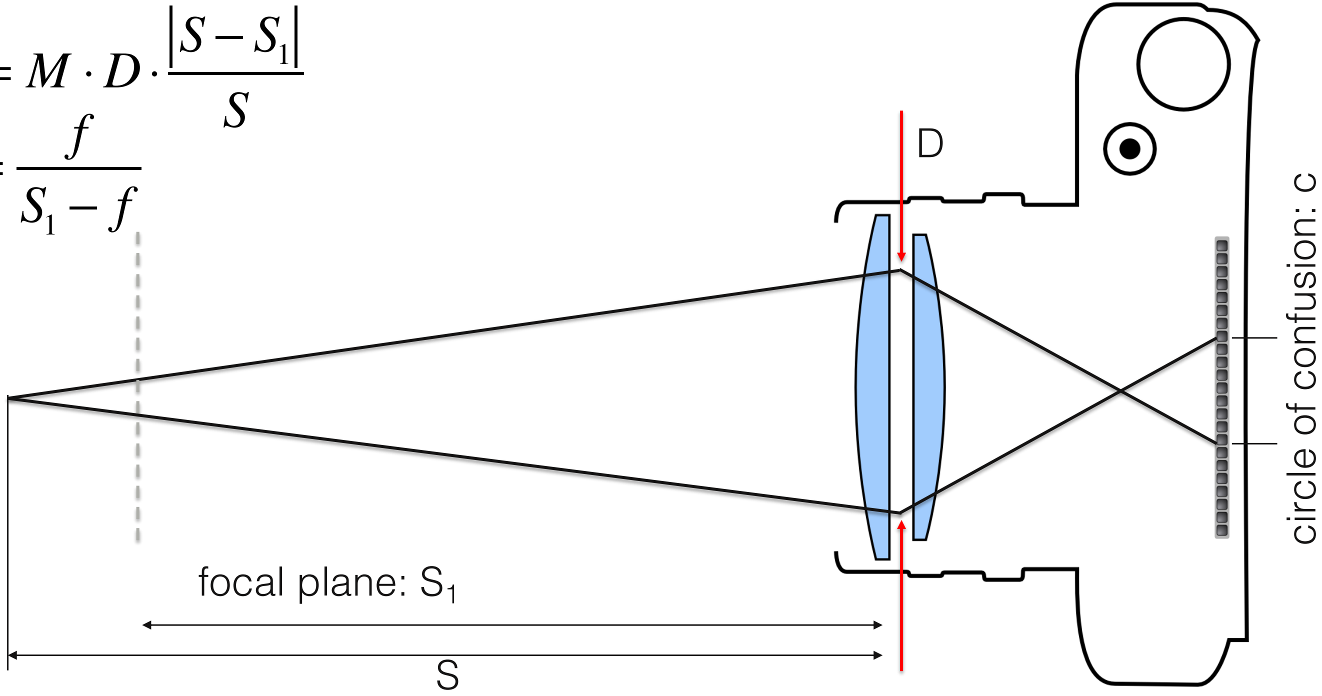

Circle of Confusion

$$c = M \cdot D \cdot \frac{|S - S_1|}{S}$$

Canon 5D Mark III: f=50mm, f/2.8 (N=2.8), focused at 5m, pixel size=7.5um

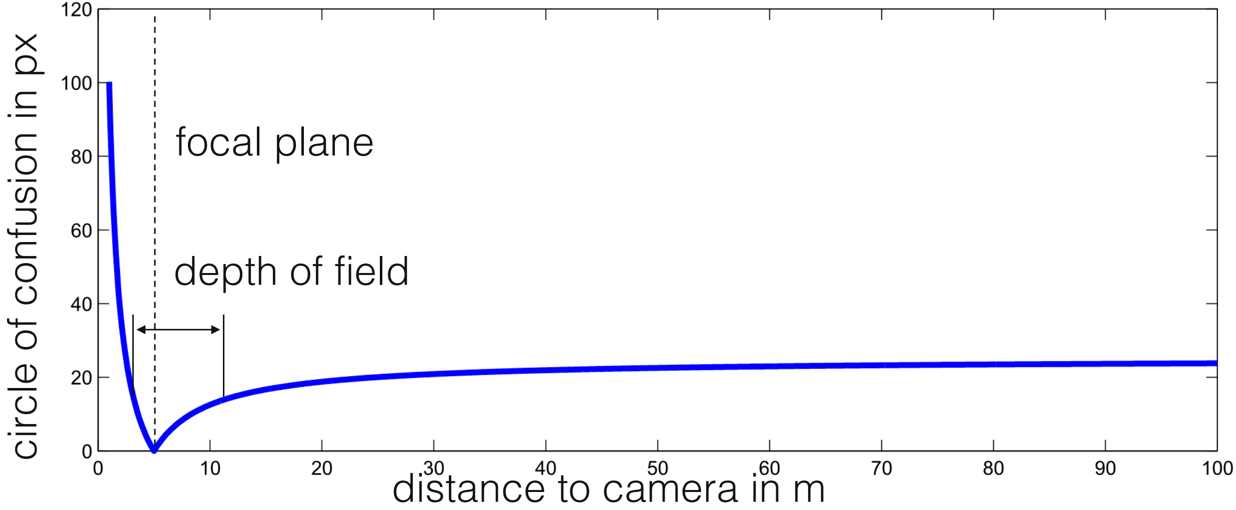

Hyperfocal Distance

$$H = \frac{f^2}{Nc}$$

<u>Canon 5D Mark III</u>: f=50mm, f/2.8 (N=2.8), focused at 5m, pixel size=7.5um

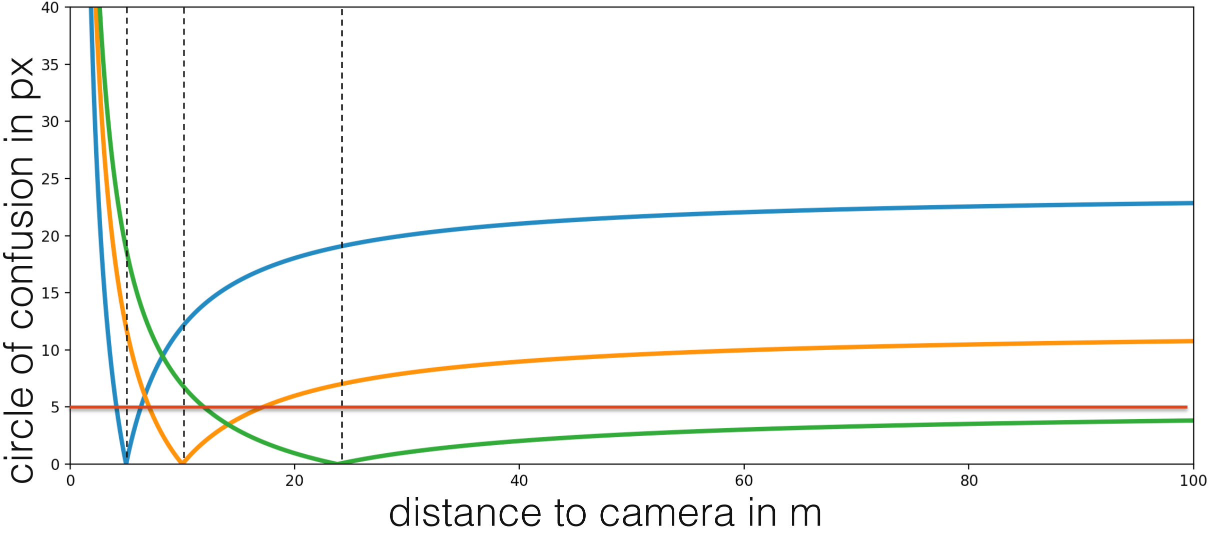

Hyperfocal Distance

$$H = \frac{f^2}{Nc}$$

Canon 5D Mark III: f=50mm, f/2.8 (N=2.8),

focused at 5m, pixel size=7.5um

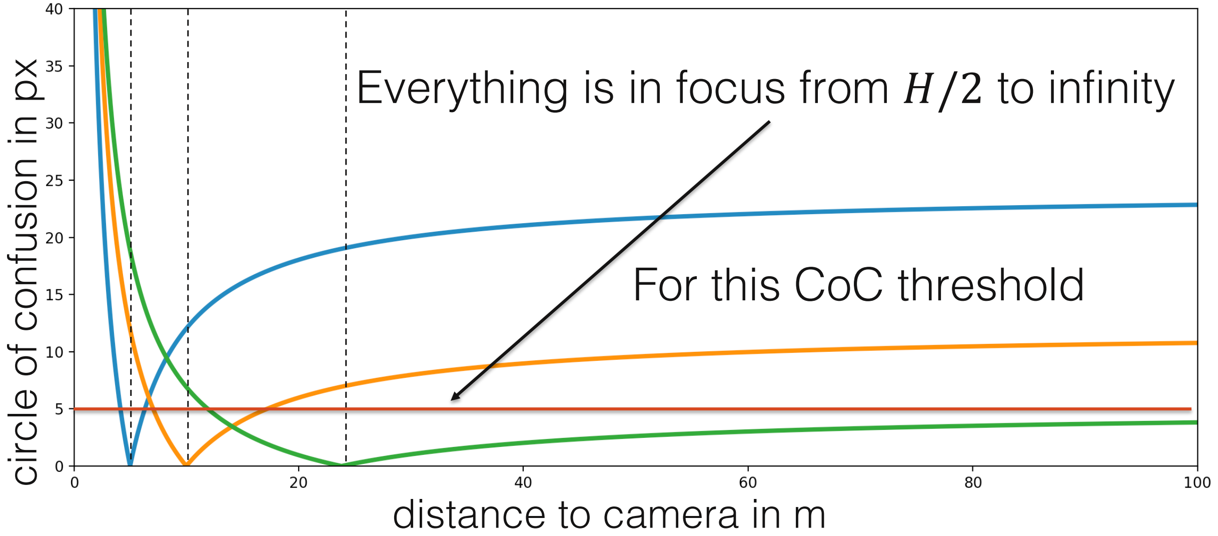

Depth of Field

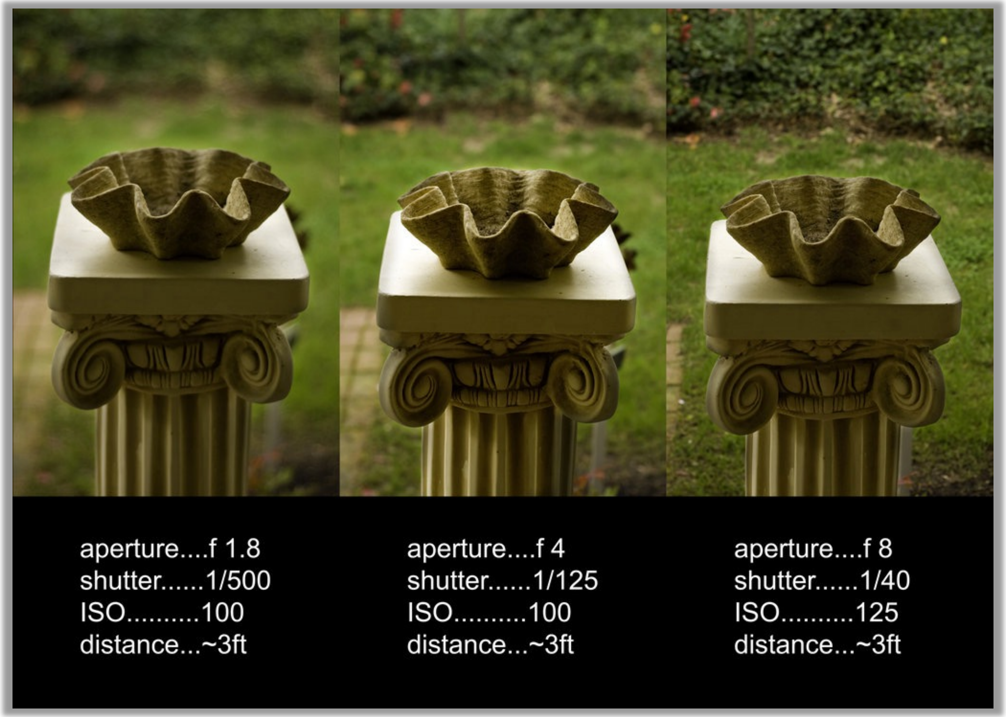

http://photographywisdom.com/

Depth of Field & Motion Blur

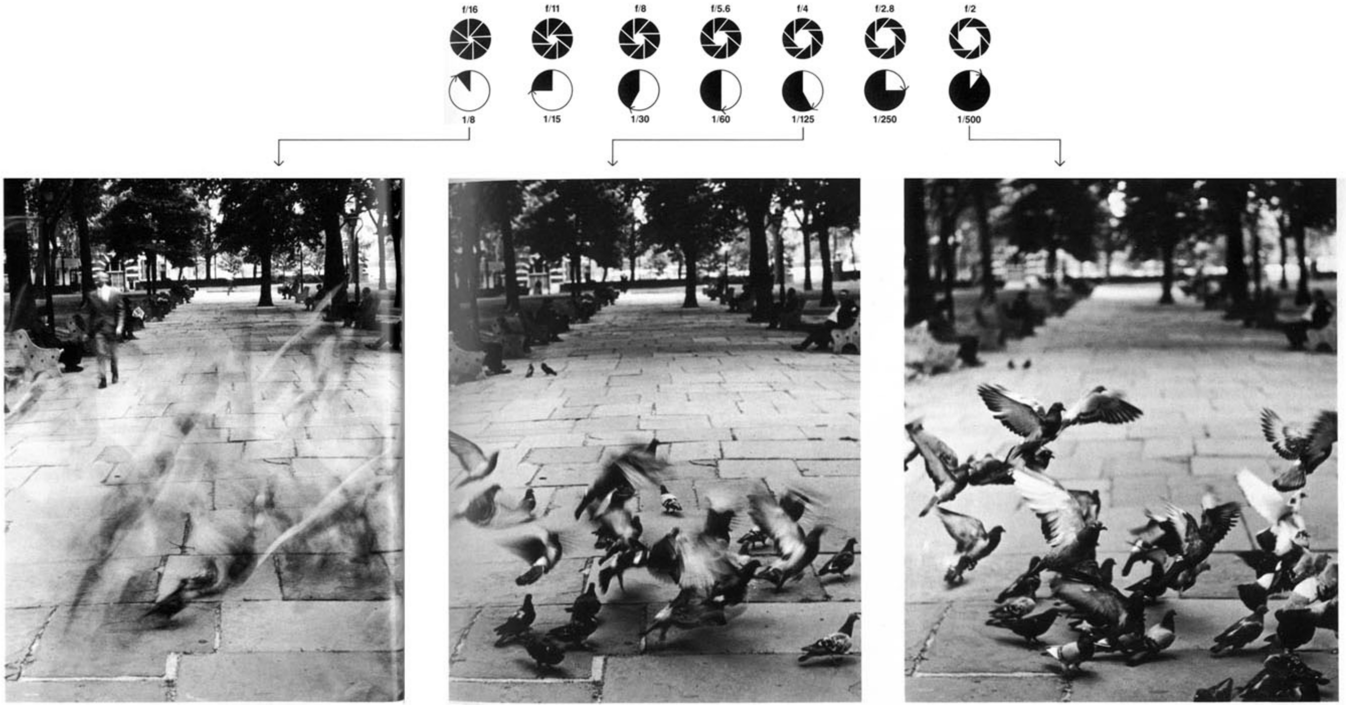

London, Photography

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

Exposure (shutter speed)

- exposure = time (e.g. 1/250, 1/60, 1, 15, bulb)

wikipedia

¼ sec, f/3.3,

2 sec, f/6.3,

Global Shutter vs. Rolling Shutter

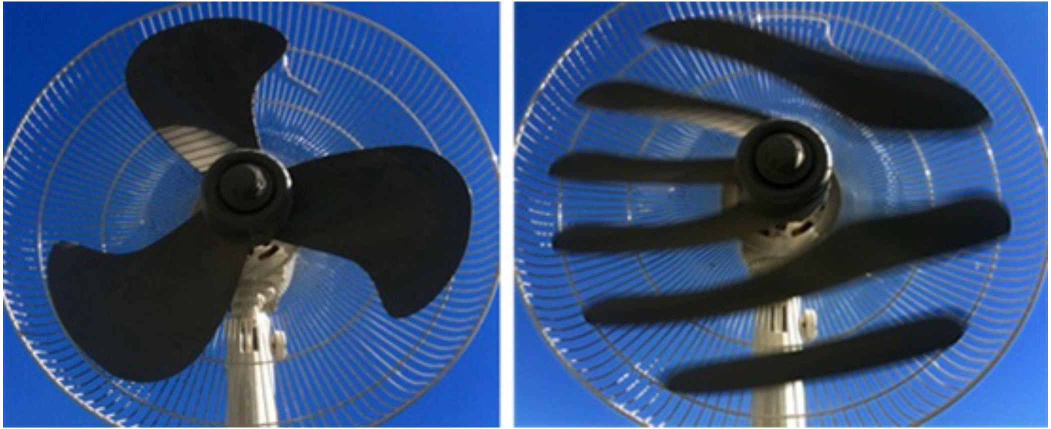

All sensor pixels exposed at same time

Row-by-row readout of image

- shorter exposure times per pixel
- motion artifacts

http://lfa.mobivap.uva.es/~fradelg/phd/notes/global-shutter.html

What are these dark bands?

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

ISO ("film speed")

sensor
sensitivity
–
analog gain
applied before
ADC!

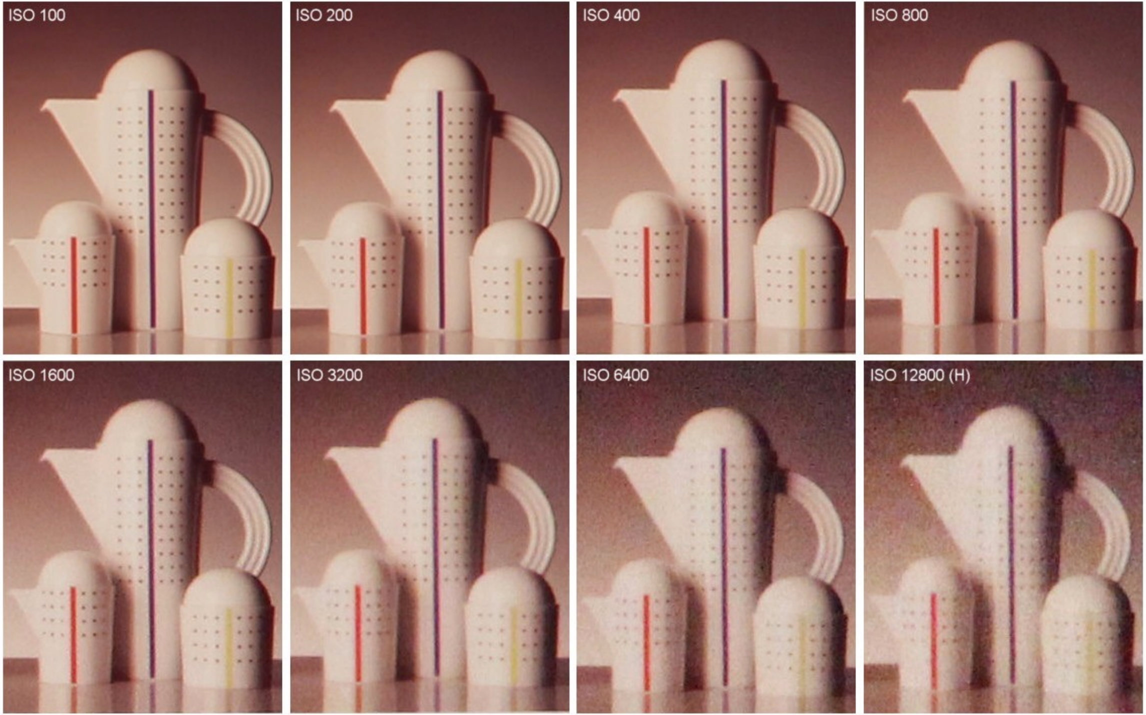

bobatkins.com

Note about the name ISO

ISO is not an acronym.

- It refers to the International Organization for Standardization.
- ISO comes from the Greek word ίσος, which means equal.
- It is pronounced (roughly) eye-zo, and should not be spelled out.

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

Dynamic Range

- ratio between largest and smallest possible value
- bit depth also important! common bit depths: 12-14 bits RAW / 8 bits JPEG

high dynamic range

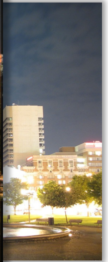

Kevin McCoy

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

What’s a Pixel?

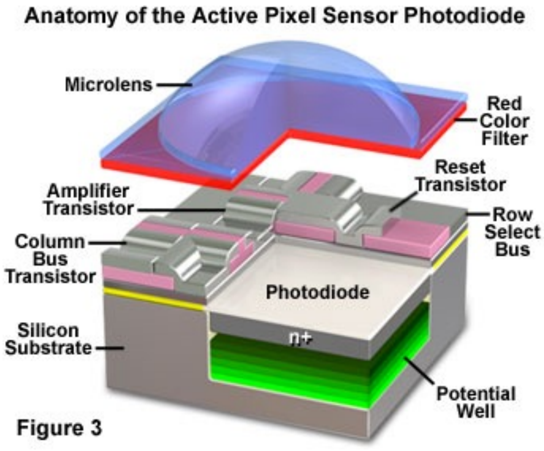

source: Molecular Expressions

photon to electron converter

→ photoelectric effect!

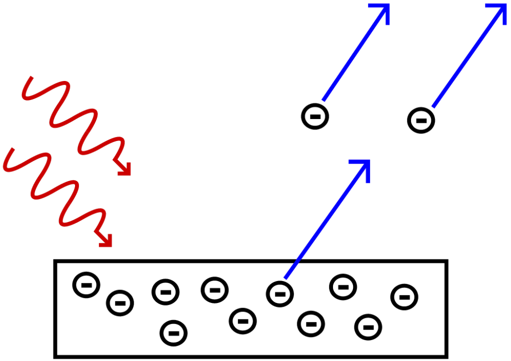

wikipedia

What's a Pixel?

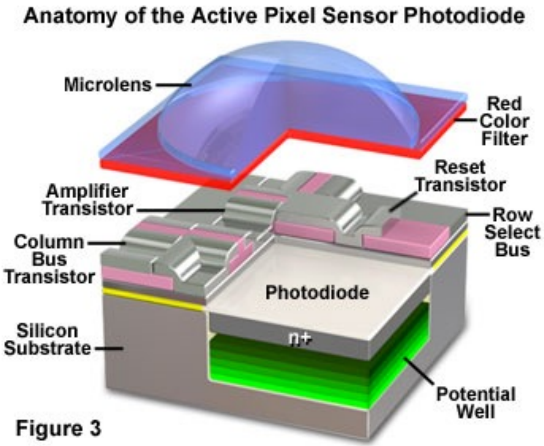

source: Molecular Expressions

- microlens: focus light on photodiode
- color filter: select color channel
- quantum efficiency: ~50%
- fill factor: fraction of surface area used for light gathering

Most Common: Color Filter Arrays

Bayer pattern

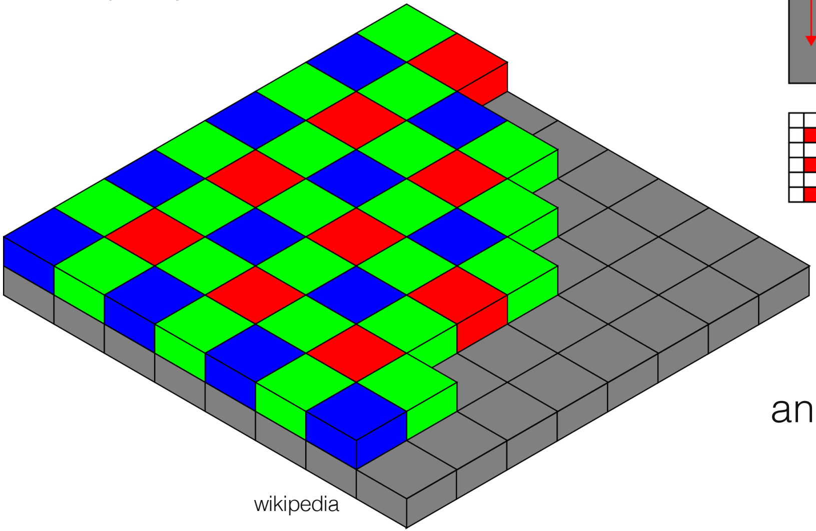

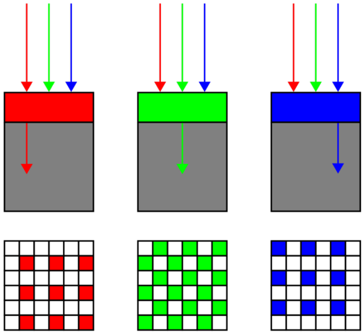

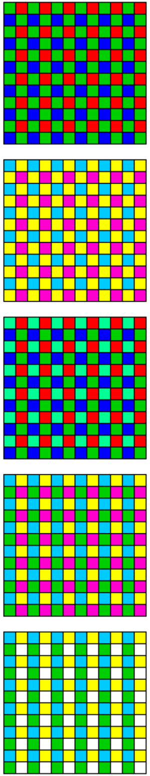

any combination possible

tradeoffs?

Digital Photography - Overview

- optics & aberrations
- field of view
- aperture
- depth of field
- exposure
- ISO
- dynamic range
- sensors
- demosaicking & denoising

camera processing pipeline

Photons to RAW Image

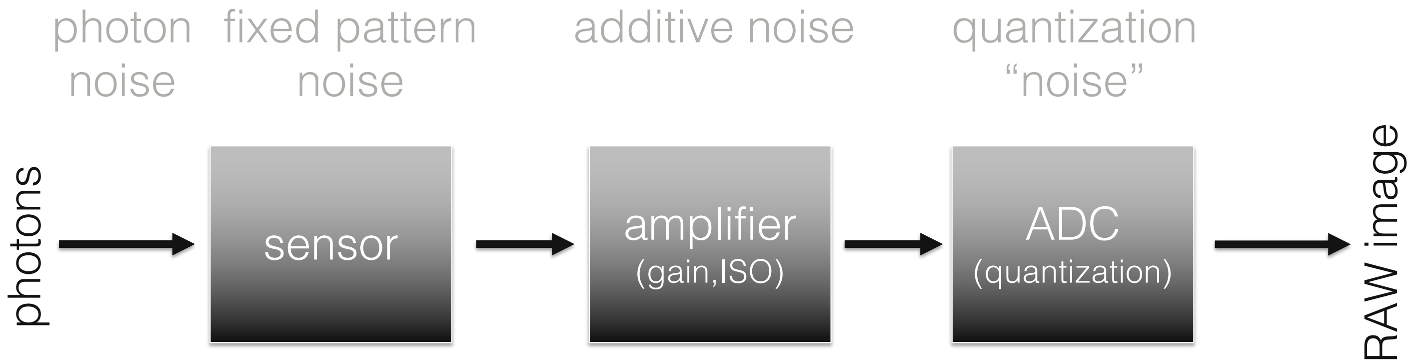

Image Processing Pipeline

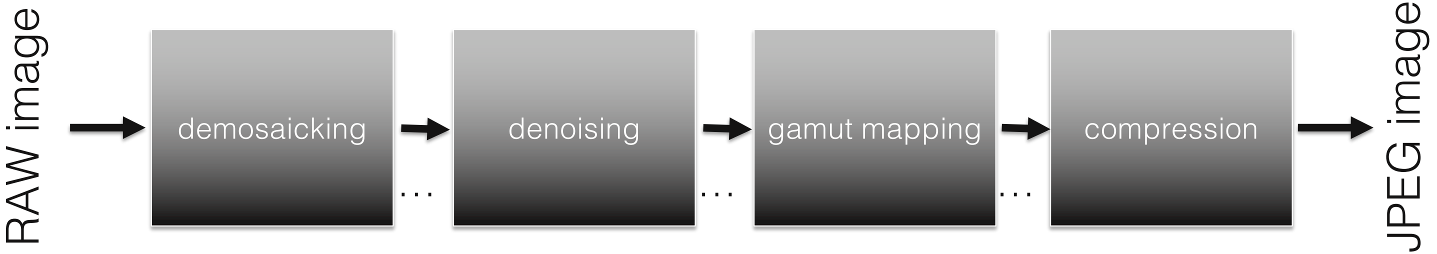

also:

- dead pixel removal
- dark frame subtraction (fixed pattern / thermal noise removal)
- lens blur / vignetting / distortion correction
- sharpening / edge enhancement

Image Processing Pipeline

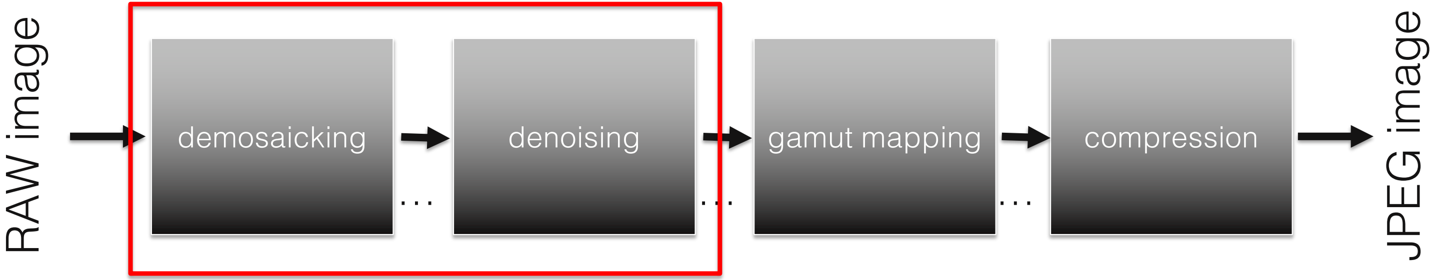

also:

- dead pixel removal
- dark frame subtraction (fixed pattern / thermal noise removal)
- lens blur / vignetting / distortion correction
- sharpening / edge enhancement

Demosaicking (CFA Interpolation)

RAW

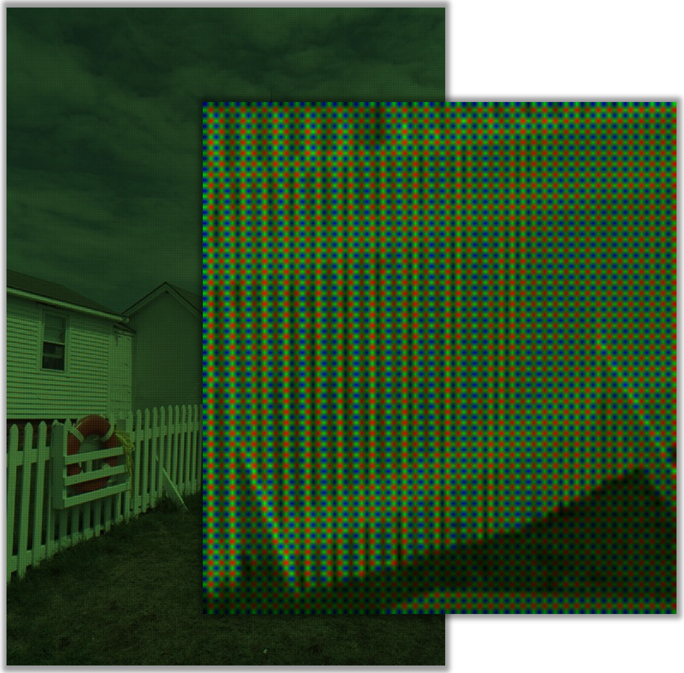

image from Kodac dataset

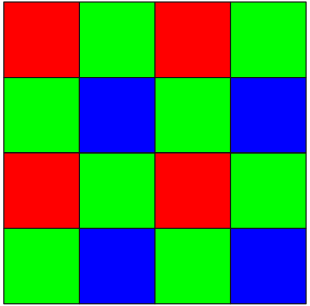

Bayer CFA

Demosaicking (CFA Interpolation)

RAW

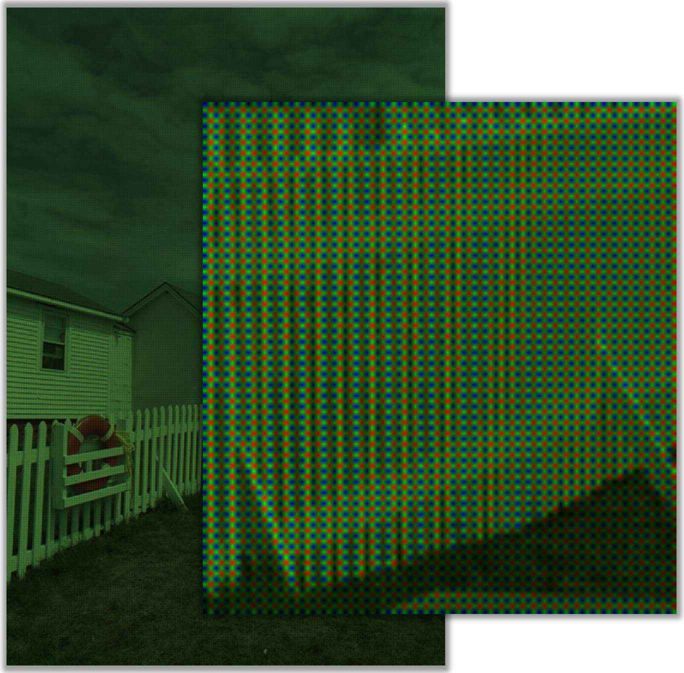

image from Kodac dataset

linear interpolation green channel

$$\hat{g}_{lin}(x,y) = \frac{1}{4}\sum_{(m,n)} g(x+m, y+n)$$

$$(m,n) = \{(0,-1),(0,1),(-1,0),(1,0)\}$$

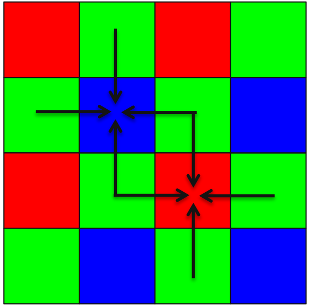

Bayer CFA

Demosaicking (CFA Interpolation)

RAW

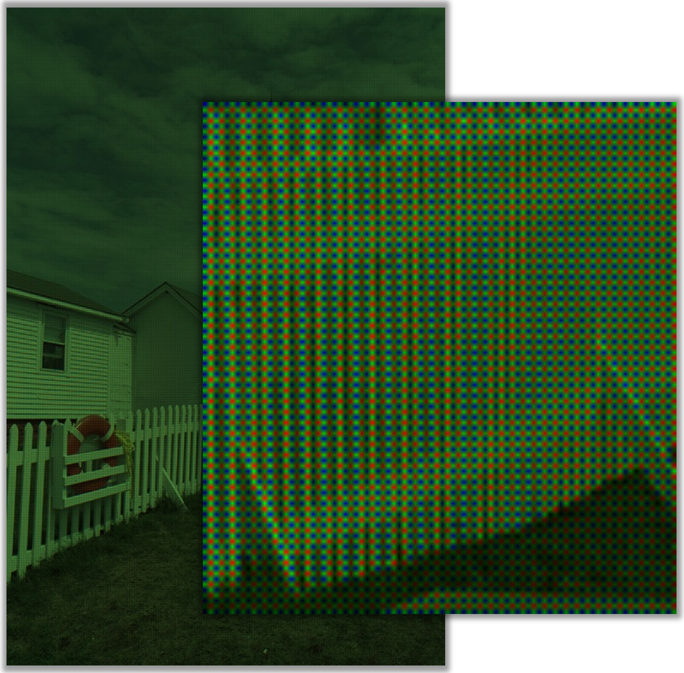

linear interpolation

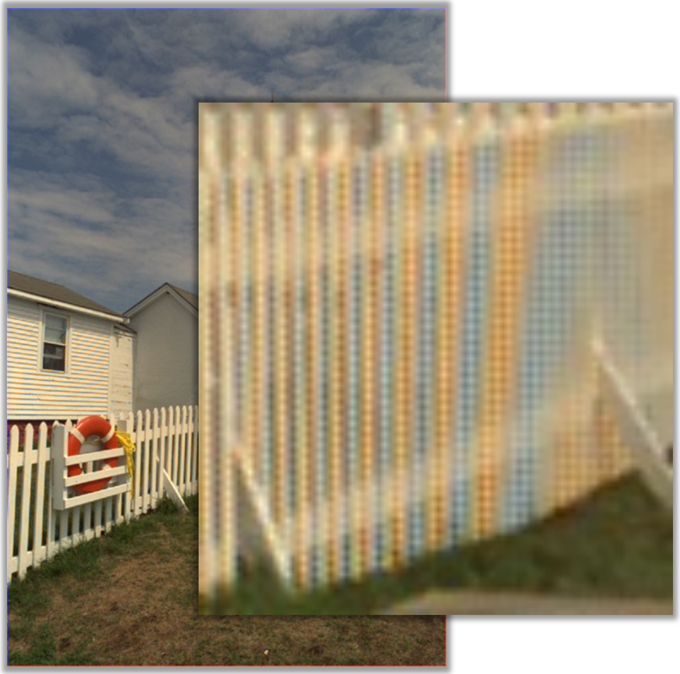

image from Kodac dataset

Demosaicking (CFA Interpolation)

original

RAW

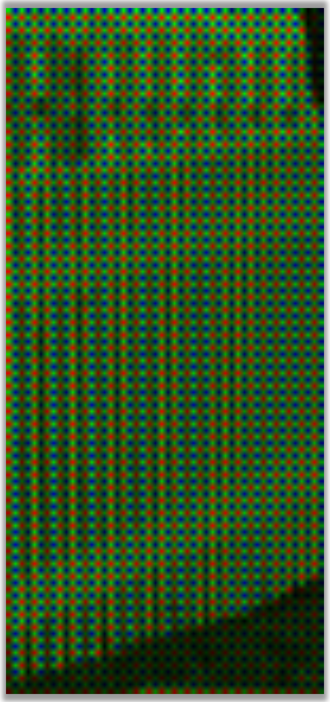

demosaicked

image from Kodac dataset

Denoising

noisy image

(Gaussian iid noise, σ=0.2)

- problem: have noisy image, want to remove noise but retain high-frequency detail

Denoising – Most General Approach

$$i_{denoised}(x) = \frac{1}{\sum_{\text{all pixels } x'} w(x,x')} \sum_{\text{all pixels } x'} i_{noisy}(x') \cdot w(x,x')$$

- many (not all) denoising techniques work like this
- idea: average a number of similar pixels to reduce noise
- question/difference in approach: how similar are two noisy pixels?

Denoising – Most General Approach

$$i_{denoised}(x) = \frac{1}{\sum_{\text{all pixels } x'} w(x,x')} \sum_{\text{all pixels } x'} i_{noisy}(x') \cdot w(x,x')$$

1. Local, linear smoothing
2. Local, nonlinear filtering
3. Anisotropic diffusion
4. Non-local methods

Denoising – 1. Local, Linear Smoothing

$$i_{denoised}(x) = \frac{1}{\sum_{\text{all pixels } x'} w(x,x')} \sum_{\text{all pixels } x'} i_{noisy}(x') \cdot w(x,x')$$

$$w(x,x') = \exp\left(-\frac{\left\| x' - x \right\|^2}{2\sigma^2} \right)$$

- naïve approach: average in local neighborhood, e.g. using a Gaussian low-pass filter

Denoising – 2. Local, Nonlinear Filtering

$$i_{denoised}(x) = median\left(W\left(i_{noisy}, x\right)\right)$$

small window of image i_{noisy} centered at x

- almost as naïve: use median filter in <u>local</u> neighborhood

Denoising

noisy image (Gaussian, σ=0.2)

Denoising – 3. Bilateral Filtering

original

Gaussian filtering

bilateral filtering

Denoising – 3. Bilateral Filtering

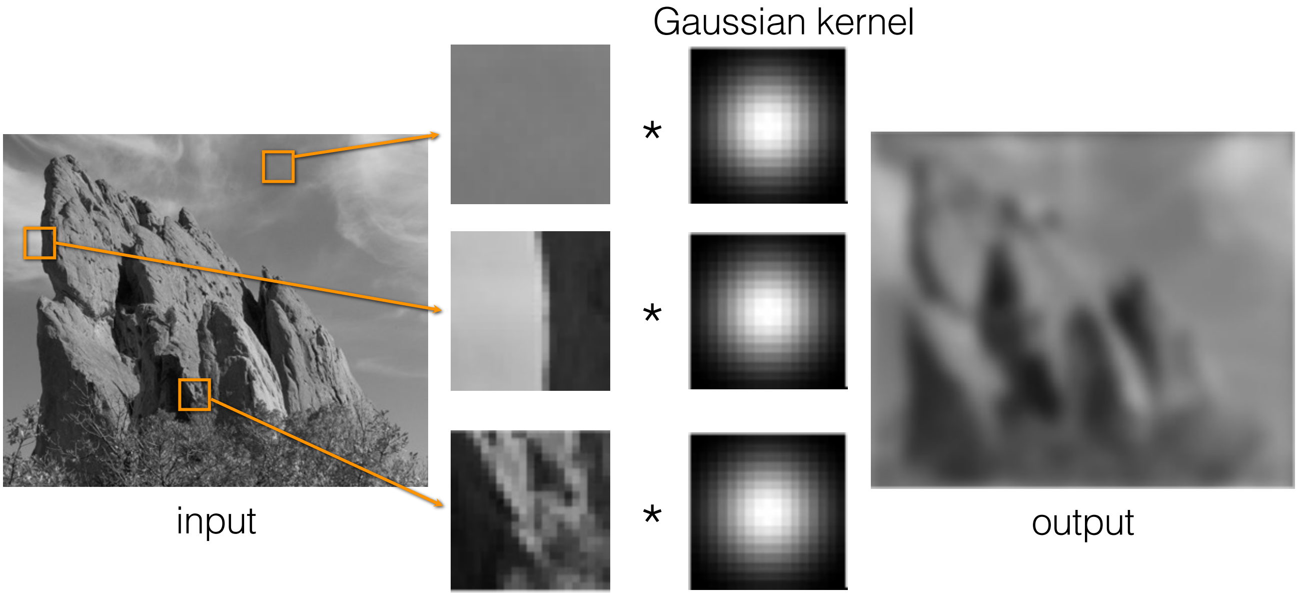

Why is the output so blurry?

Denoising – 3. Bilateral Filtering

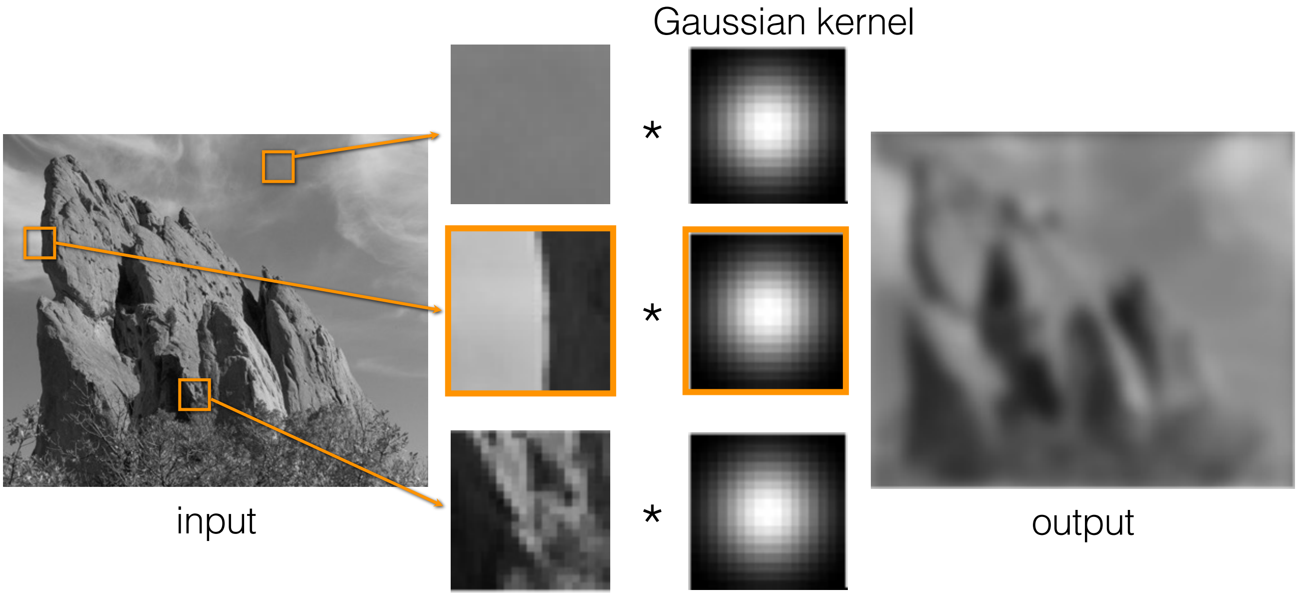

Blur kernel averages across edges

Denoising – 3. Bilateral Filtering

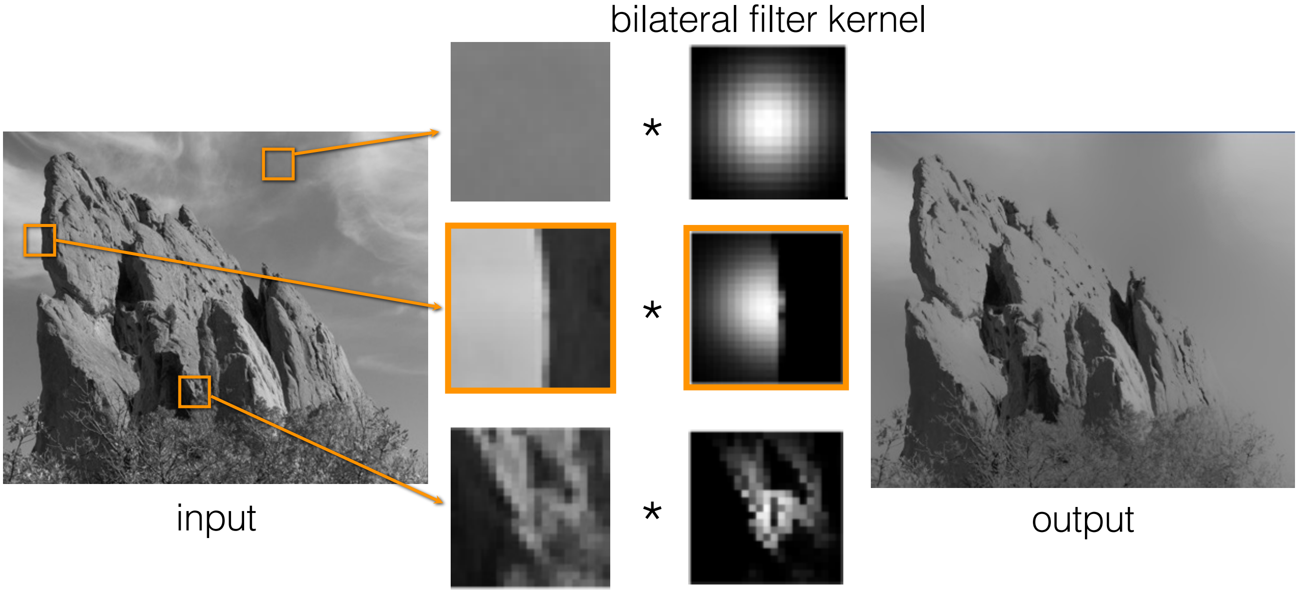

Do not blur if there is an edge! How does it do that?

Denoising – 3. Bilateral Filtering

$$i_{denoised}(x) = \frac{1}{\sum_{\text{all pixels } x'} w(x,x')} \sum_{\text{all pixels } x'} i_{noisy}(x') \cdot w(x,x')$$

spatial distance

distance of intensities

$$w(x,x') = \boxed{\exp\left(-\frac{\left\| x' - x \right\|^2}{2\sigma^2}\right)} \cdot \boxed{\exp\left(-\frac{\left\| i_{noisy}(x') - i_{noisy}(x) \right\|^2}{2\sigma_i^2}\right)}$$

- more clever: average in <u>local</u> neighborhood, but only average similar intensities!

Denoising – Gaussian Filter

J: filtered output (is blurred)
f: Gaussian convolution kernel
I: step function & noise

$$J(x) = \sum_{\xi} f(x,\xi) \quad I(\xi)$$

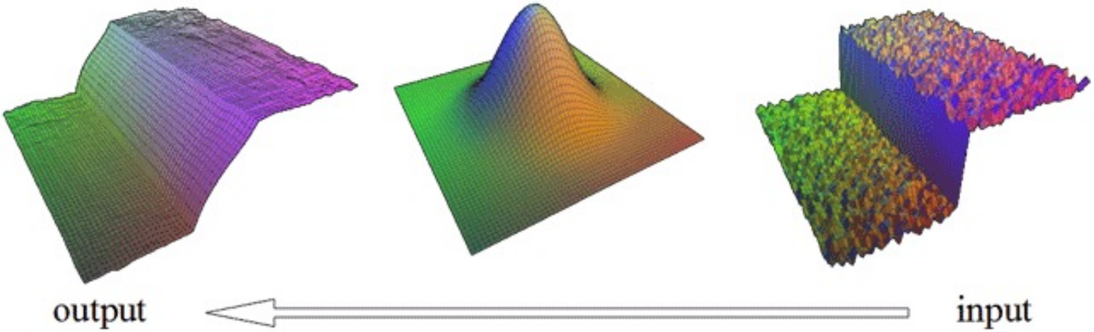

Denoising – Bilateral Filter

J: filtered output (is not blurred)
f: Gaussian convolution kernel
I: noisy image (step function & noise)

difference in intensity as scale!

$$J(x) = \frac{1}{k(x)} \sum_{\xi} f(x,\xi) \quad g(I(\xi)-I(x)) \quad I(\xi)$$

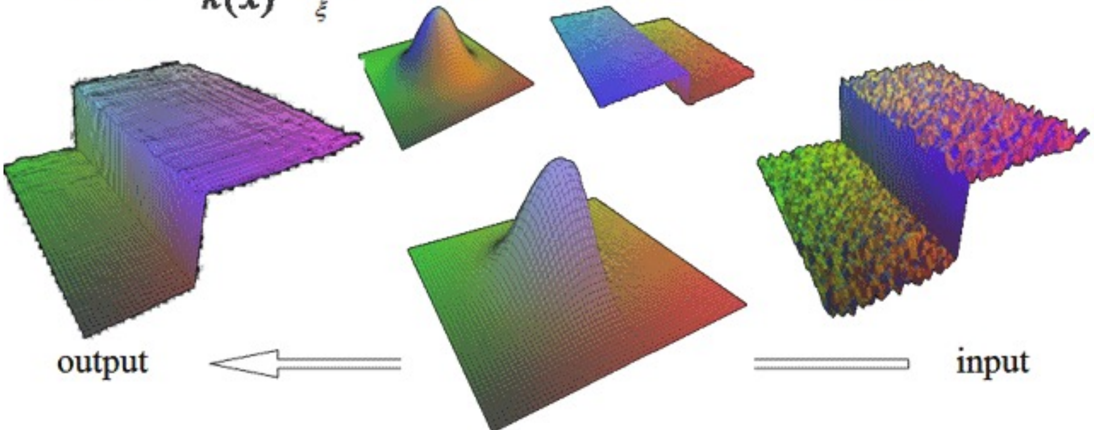

Tomasi & Manduchi, 1998

Denoising – Bilateral Filter

original image

bilateral filter = “edge-aware smoothing”

Denoising – Bilateral Filter

noisy image

bilateral filter = “edge-aware smoothing”

Exploring the bilateral filter parameter space

input

$\sigma_r = 0.1$ | $\sigma_r = 0.25$ | $\sigma_r = 8$ (Gaussian blur)

$\sigma_s = 2$

$\sigma_s = 6$

$\sigma_s = 18$

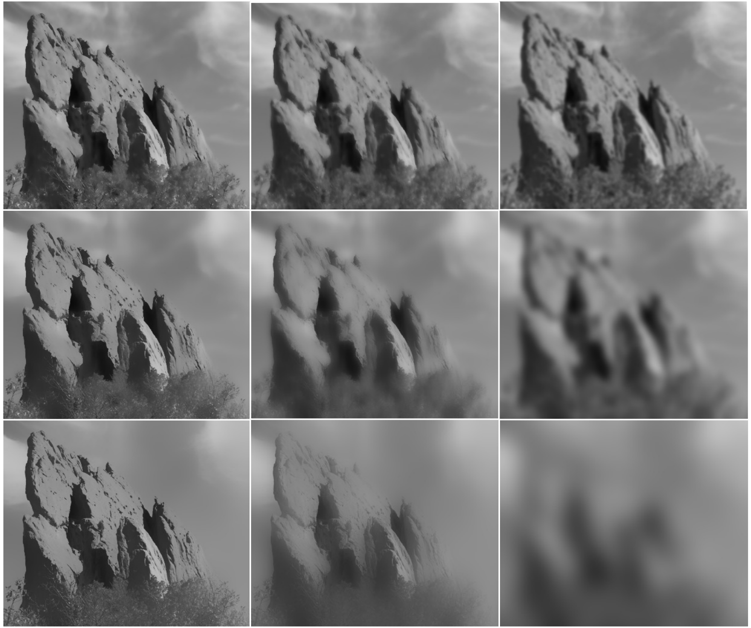

Denoising

noisy input

bilateral filtering

median filtering

Contrast enhancement

How would you use Gaussian or bilateral filtering for sharpening?

input

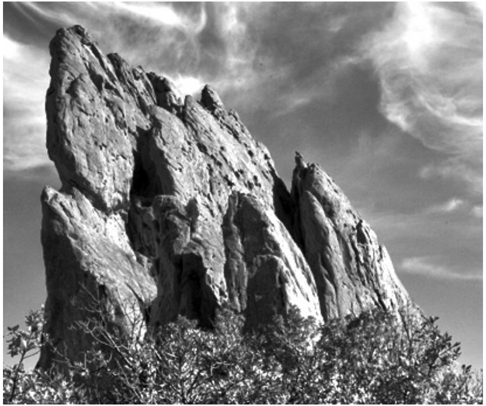

sharpening based on bilateral filtering

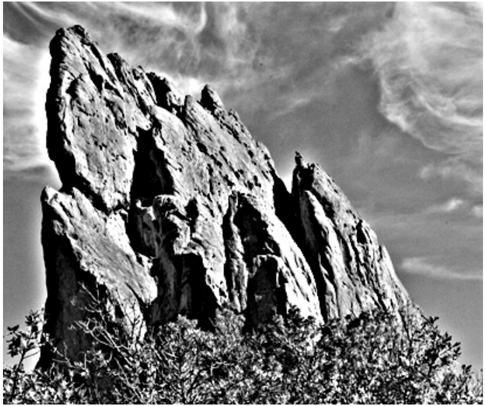

sharpening based on Gaussian filtering

Photo retouching

Photo retouching

original

digital pore removal (aka bilateral filtering)

Before

After

Close-up comparison

original

digital pore removal (aka bilateral filtering)

Cartoonization

input

cartoon rendition

Cartoonization

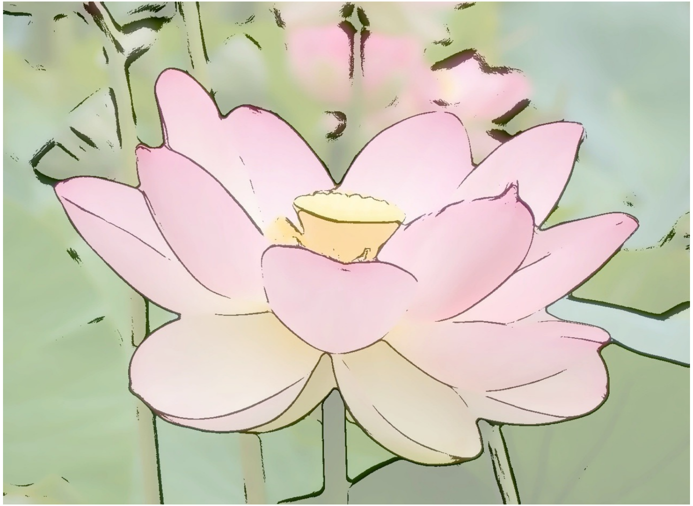

How would you create this effect?

Cartoonization

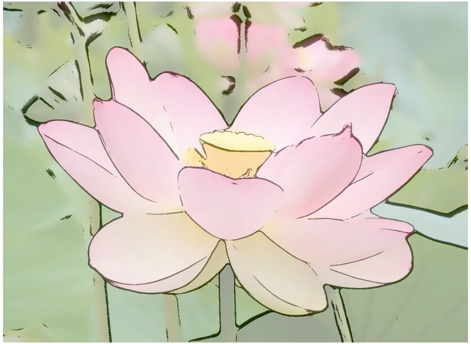

edges from bilaterally filtered image

+

bilaterally filtered image

=

cartoon rendition

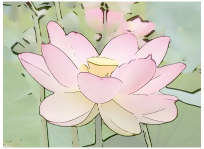

Note: image cartoonization and abstraction are very active research areas.

Denoising – 4. Non-local Means

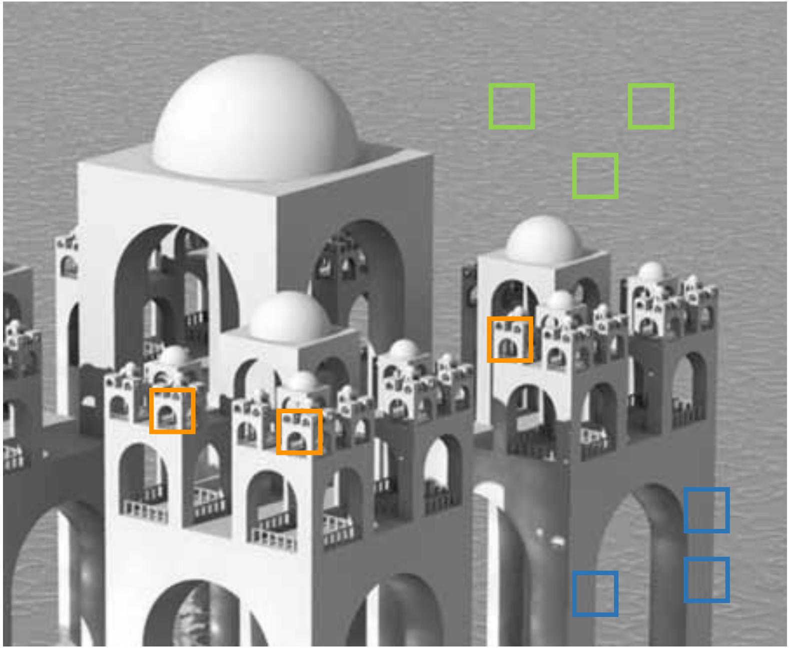

Denoising – 4. Non-local Means

- define distance between global image patches
- average distant pixels with similar neighborhood!

$$i_{denoised}(x) = \sum_{\text{all pixels } x'} i_{noisy}(x') \cdot w(x, x')$$

[Buades 2005]

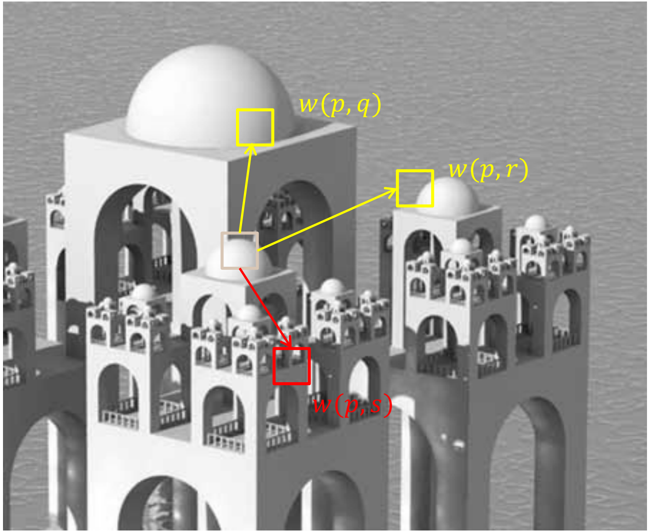

Denoising – 4. Non-local Means

$$i_{denoised}(x) = \frac{1}{\sum_{\text{all pixels } x'} w(x,x')} \sum_{\text{all pixels } x'} i_{noisy}(x') \cdot w(x,x')$$

$$w(x,x') = \exp\left(-\frac{\left\| W\left(i_{noisy}, x'\right) - W\left(i_{noisy}, x\right) \right\|^2}{2\sigma^2} \right)$$

[Buades 2005]

- very powerful approach: exploit self-similarity in image; average pixels wit a similar neighborhood, but don't need to be close → non-local

Denoising – 4. Non-local Means

noisy GT Gaussian filtering

median filter bilateral filtering NL-means

Denoising – 4. Non-local Means

noisy GT Gaussian filtering

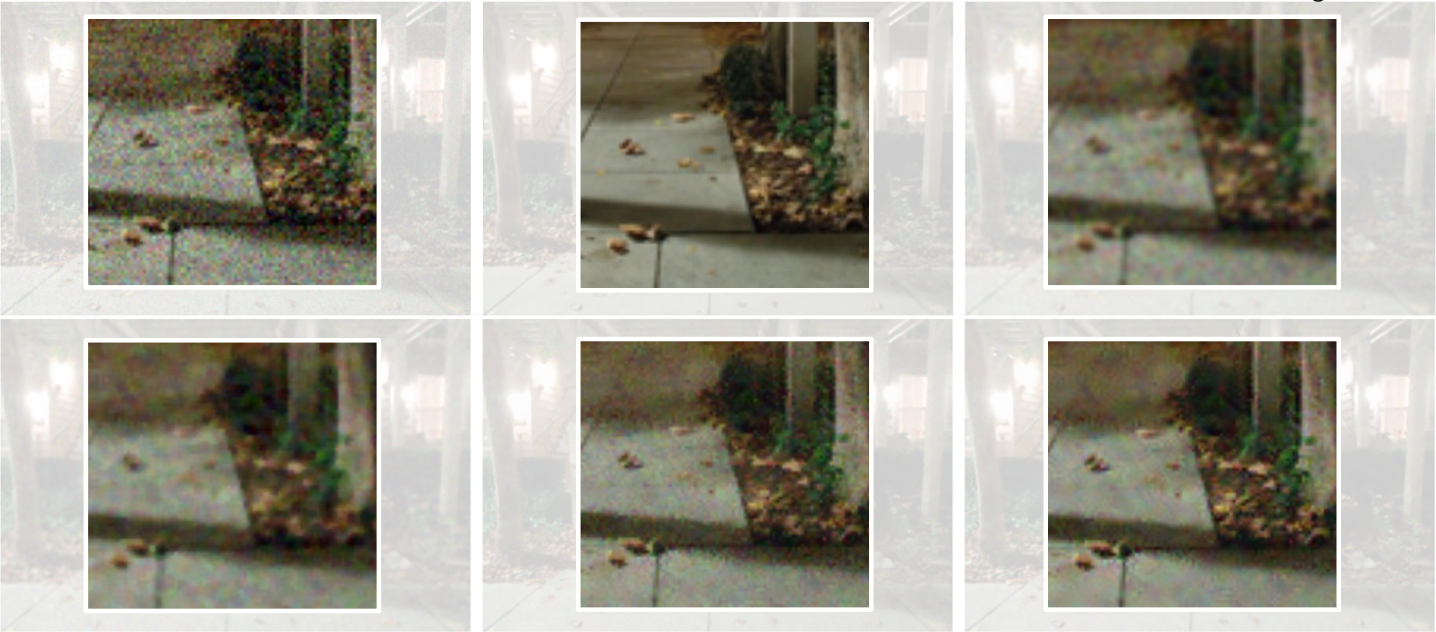

median filter bilateral filtering NL-means

Everything put together

Gaussian filtering

Smooths everything nearby (even edges)
Only depends on spatial distance

Bilateral filtering

Smooths ‘close’ pixels in space and intensity
Depends on spatial and intensity distance

Non-local means

Smooths similar patches no matter how far away
Only depends on intensity distance

Next: Radiometry & Photometric Cues

References and Further Reading

- London, Upton, Stone, “Photography”, Pearson, 11th edition, 2013
- Stanford CS 178, “Digital Photography”, Course Notes
- CMU Computational Photography course
- wikipedia